

Optimal Windows for Aggregating Ratings in Electronic Marketplaces

Christina Aperjis

Department of Management Science and Engineering, Stanford University, caperjis@stanford.edu

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 the Window Aggregation Mechanism, a widely used class of mechanisms that shows the average value of the seller’s ratings within some fixed window of past transactions. We suggest approaches for choosing the window size 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

1. Introduction

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).

If there is a *reputation mechanism* in place, dishonesty involves a greater immediate payoff at the expense of a lower long-term payoff. In such a mechanism, after a transaction the buyer can rate the seller; all past ratings are aggregated into the seller’s score. When a seller posts a description

for an item, potential buyers also observe the seller's score. 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). In our model the expected payment to the seller is thus a function of his description and his score, and may also depend on the mechanism by which ratings are aggregated. Empirical studies have shown that sellers with high scores enjoy a *price premium*: on average they sell at higher prices than less reputable sellers; see Resnick et al. (2006) for a survey. Motivated by these studies, we assume that the expected payment to the seller is increasing in his score.

A natural goal for a reputation mechanism is to encourage sellers to describe their items truthfully. While incentivizing truthfulness is intuitively appealing, it is also connected to *efficient trade*. Trade is efficient if the item is allocated to the agent that values it most. The buyers' valuations depend on the item's characteristics, which only the seller knows before an online sale. Thus, it is important that the seller gives an accurate description to potential buyers.

We study the *Window Aggregation Mechanism*, a widely used mechanism in which the seller's score is the average value of his T most recent ratings. With the goal of incentivizing the seller to describe the item truthfully, we address the *design question* of choosing the right window size T . We define an optimal window size as one which maximizes the range of parameters for which it is optimal for the seller to be always truthful, and we study the dependence of the optimal window size on the parameters of the model.

Our main result is an interesting qualitative tradeoff: informally, *increasing the window size is more likely to make patient sellers truthful, while it is less likely to make high quality sellers truthful*. First, increasing the window size is more likely to make patient sellers truthful. This is an intuitive result, since *patient sellers value future payments more, and thus an aggregation mechanism with longer memory can successfully couple current behavior with distant future payoffs*. On the other hand, a larger window is less likely to make high quality sellers truthful. When the seller has high value items for sale most of the time 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

have high value. 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.*

Alternatively, we can interpret our main result by directly looking at the effect of the window size on the seller's future scores. By choosing a window size T , the system designer determines *how much* and for *how long* the seller's future scores decrease if he does not describe his current item truthfully. We identify the following tradeoff between the intensity and the duration of this score reduction: the intensity is decreasing in T , while the duration is increasing in T . The duration effect is more important for patient sellers, so the optimal window size increases. On the other hand, the intensity effect becomes more important for high quality sellers, so the optimal window size decreases.

In Section 3 we introduce our design framework, which we use in subsequent sections to show the tradeoff between incentivizing patient and high quality sellers for a range of settings. First, we consider the case of *perfect monitoring* (Section 4), where the buyer rates the seller accurately after every transaction. Then, we consider two types of *imperfect monitoring*: the seller may not receive ratings after some transactions (Section 5), and the rating that the seller receives may not accurately reflect his action (Section 6). In Section 7 we discuss the mapping between a fixed number of transactions and a fixed interval of time. We conclude in Section 8 by relating our findings to management practice. The appendix contains proofs of the results in the paper.

In our model, the objective of the aggregation mechanism is to induce sellers to behave in a way that benefits the market. 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* behavior of the buyers; that is, a fixed payment function which only depends on the information available to buyers (i.e., the window size, the seller's score, and the description of the item). We believe this nonequilibrium approach may be reasonable in practice because the large and dynamic set of participants in the major online markets makes the rationality, knowledge, and

coordination required for equilibrium difficult to ensure. At the very least, it seems reasonable that short run behavior in these markets may not be an equilibrium. Along these lines, Fan et al. (2005) also take a non-equilibrium approach and assume that the seller has a belief over the average bidder behavior; however, they consider a distinct class of mechanisms from our work, where exponential smoothing is employed.

In this paper, we study the Window Aggregation Mechanism, and prove results on how the optimal window size is affected by the seller's attributes. Although Window Aggregation Mechanisms are widely used in electronic marketplaces because of their simplicity, there are other possible ways to aggregate ratings. For example, instead of using a finite window size, ratings might be averaged over the entire lifetime of the seller; this approach has been shown to be ineffective in a number of settings (e.g., Fan et al. 2005, Aperjis and Johari 2008, Cripps et al. 2004). It is possible to design good aggregation mechanisms by averaging over the lifetime of the seller if different weights are given to ratings from different periods. The work of Fan et al. (2005) is one such example: they propose exponential smoothing as a good way to aggregate ratings, and evaluate it through simulations. In the same vein, for a more general class of Weighted Reputation Mechanisms, we establish (under a related but different set of assumptions) the same tradeoff identified above: averaging over a longer past history of ratings is *more likely* to incentivize patient sellers to be truthful, but *less likely* to incentivize high quality sellers to be truthful (Aperjis and Johari 2008). Despite these properties, such mechanisms are not characterized by the simplicity or ubiquity of Window Aggregation Mechanisms.

2. 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 buyers' behavior and the market mechanism at each time period; instead, we 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 a random variable X which takes values in $[v_L, v_H]$, where $0 \leq v_L < v_H \leq 1$. In particular, we denote by v_L (resp., v_H) the minimum

(resp., maximum) value of an item that the seller may have for sale. We assume that the values of items in different periods are independent and identically distributed, and define $q \equiv EX$ to be the expected value of the items for sale by the seller. The seller observes the value of the item at the beginning of a period and decides what description to post. Potential buyers observe the seller's description and the average value of the ratings that the seller received in the last T periods. The expected payment to the seller is a function of this aggregate information and the seller's description.

Asymmetric information may create two kinds of problems in online markets. First, sellers may not deliver on their promised action (moral hazard). Second, sellers may have hidden types that are not known by the mechanism designer (adverse selection). Qualitatively, our model has elements of both moral hazard and adverse selection: a seller may not describe the item he has for sale accurately, and the mechanism designer may not know certain characteristics of sellers (such as patience and quality). Models for online marketplaces that combine moral hazard and adverse selection have also been considered by Cabral and Hortacsu (2009) and Li (2009).

After purchase, the buyer rates the seller. We assume that the rating depends (either deterministically or randomly) on *how much the seller exaggerates the value of the item in his description*. Specific models for this dependence are introduced in subsequent sections. We note that other factors may also affect the rating; however, since our goal is to incentivize the seller to describe the item truthfully, we focus on the dependence of the rating on how much the seller exaggerates the item's value.

Let r_i be the value of the rating that the seller received i periods ago, and let $\vec{r} = (r_1, r_2, \dots)$ be the vector of ratings that the seller has received up to now. As mentioned above, we assume that potential buyers have access to the mean value of the ratings that the seller received in the last T periods, for some $T \geq 1$. In particular, potential buyers see

$$\frac{1}{T} \sum_{i=1}^T r_i,$$

which we call the seller's *score*. To compute this score, the aggregation mechanism needs to keep information on the T most recent ratings of the seller. Both the seller and the mechanism have

access to \vec{r} , but potential buyers only observe the seller's score. We call this mechanism the *Window Aggregation Mechanism*, and we refer to T as the *window size*.

The Window Aggregation Mechanism is widely used in online marketplaces, such as eBay and the Amazon Marketplace. In both marketplaces the average value of the seller's ratings in the last twelve months is shown in a prominent position: next to the description of the item on eBay and below the seller's name on Amazon. Additional information, such as the seller's scores for other window sizes,¹ is available in subsequent pages. Cabral and Hortacsu (2009) conduct an empirical study that shows that the information that is prominently shown to buyers has a larger effect on the expected payment to the seller. This motivates us to study the use of a single window size.

It is worth noting that online markets use a fixed interval of *time* rather than a fixed number of transactions. To avoid this potential complexity, for the moment we simply assume that the seller has an item for sale in every period. More generally, we can consider a model where the seller has multiple items for sale in each period; using such a model, we further discuss the mapping between a fixed number of transactions and a fixed interval of time in Section 7.

We assume that the expected payment to the seller is $v \cdot b_T(s)$, where s is his score and v the description he posted.² The expected payment to the seller is thus increasing in the description he posts, as is reasonable to assume in practice. We note that if $v_L > 0$, then the seller receives a premium even for the lowest possible value items. We call $b_T(s)$ the *premium function*. This function measures how much buyers are willing to pay to transact with sellers that have higher scores. In particular, when $s_1 < s_2$ and the seller advertises an item of value v , buyers are willing to pay $v \cdot (b_T(s_2) - b_T(s_1))$ more on average to buy the item when the seller has score s_2 instead of s_1 . Motivated by empirical studies (e.g., Ghose et al. 2005, Cabral and Hortacsu 2009), we assume that the premium function is increasing in the seller's score. We use the subscript T on the premium function to denote that the buyers may react to the window size that is being used.

¹The Amazon Marketplace shows the mean value of ratings out of five stars in the last twelve months, as well as the percentage of positive, neutral and negative feedback that the seller received within the last 30, 90, and 365 days. On the other hand, eBay shows the absolute numbers of positive, neutral and negative feedback for the past 1, 6, and 12 months.

²In this paper we assume that the expected payment to the seller depends on the seller's description, the window size T , and the average value of the seller's T most recent ratings. In general, the expected payment to the seller may also depend on other statistics—if such statistics are provided by the marketplace. However, in this paper we do not consider this possibility and focus on a single aggregate statistic.

The seller chooses a policy that is a best response to the premium function $b_T(\cdot)$. In our model, we emphasize that the seller is not intrinsically honest or dishonest; he is rational and chooses the description that maximizes his payoff. We assume that the seller chooses a description in $[v_L, v_H]$; however, our results hold even if the maximum possible description is a constant that is greater than v_H .

We conclude with the following assumptions; *these assumptions remain in force throughout the paper.*

ASSUMPTION 1. *There exists some $T \geq 1$, and $q, \delta < 1$ such that a seller with attributes q, δ is always incentivized to be truthful under a Window Aggregation Mechanism with window size T .*

ASSUMPTION 2. *The premium function $b_T(\cdot)$ is logarithmically concave for all T .*

Assumption 1 says that the premium function is such that there exists a T at which a seller with suitable q and δ is incentivized to be always truthful; an assumption that makes the optimization problem considered in the following sections meaningful. In the case of perfect monitoring, this assumption holds for a large class of premium functions, which includes strictly convex functions. It also holds for concave functions as long as the minimum value that the seller may have for sale, i.e., v_L , is sufficiently greater than zero (Aperjis and Johari 2008). When monitoring is imperfect, Assumption 1 still holds for a large class of premium functions as long as inaccurate and missing ratings are not too frequent.

We next comment on Assumption 2. A function is logarithmically concave if its logarithm is concave. All concave functions are logarithmically concave. Other examples of logarithmically concave functions are x^n and e^{nx} for all $n \geq 1$, and the logistic function $1/(1 + e^{a-bx})$, where $b > 0$.

We note that various empirical studies suggest that the expected premium is logarithmically concave in the mean value of ratings of the seller (e.g., Cabral and Hortacsu 2009, Lucking-Reiley et al. 2007, Ghose et al. 2005). These studies regress the payment to the seller or its logarithm against some function of the average rating of the seller or some other function of the number of positive and negative ratings (in the case of eBay). Even though these studies consider all the ratings that the seller has received in his lifetime or in the last twelve months, we can get

some insight in the dependence of the payment on the percentage of positive ratings in the last T transactions by fixing the total number of transactions to T .

Cabral and Hortacsu (2009) use a data set from eBay and regress the logarithm of the price against the percentage of positive ratings of the seller. Since the logarithm of the price is a linear function of the average rating, the expected payment is logarithmically concave in the seller's score (by definition). Ghose et al. (2005) use a data set from the Amazon Marketplace and regress the logarithm of the premium against the mean rating of the seller in the last twelve months. We note that the Amazon Marketplace asks buyers to rate sellers out of five stars. Again, this corresponds to a logarithmically concave function. Lucking-Reiley et al. (2007) use a data set from eBay and regress the price against the logarithm of the number of positive ratings and the logarithm of the number of negative ratings. Setting T equal to the sum of the number of positive and negative ratings, we observe that according to this regression the price is a logarithmically concave function of the seller's score.

Key notation from this and subsequent sections is summarized in Table 1.

3. Design Framework

This paper considers the problem of *designing* a good aggregation mechanism; in particular, we consider a setting where the system designer wishes to choose the optimal window size in a Window Aggregation Mechanism. We assume that the mechanism designer's goal is to choose a window size that *maximizes the range of seller parameters q and δ for which truthfulness can be guaranteed*. Ultimately, our analysis of this design problem lends qualitative insight into the design of aggregation mechanisms: we find that averaging over a longer past history of ratings is *more likely* to incentivize patient sellers to be truthful, but *less likely* to incentivize high quality sellers to be truthful.

In this section we introduce a general design framework, which we use in subsequent sections to show the tradeoff between incentivizing patient and high quality sellers for a range of settings. In Section 3.1 we show the main result of the design framework: if two key properties are satisfied, then the optimal window is increasing in δ and decreasing in q (Theorem 1). In Section 3.2 we map out our application of Theorem 1 in subsequent sections where we consider perfect monitoring

Notation	Definition	Introduced in
δ	discount factor	Section 2
v_L	lowest possible value of each item of the seller	Section 2
v_H	highest possible value of each item of the seller	Section 2
q	expected value of each item of the seller	Section 2
T	window size	Section 2
r_i	rating seller received i periods ago	Section 2
\vec{r}	vector of all past ratings	Section 2
s	seller's score: mean value of ratings in last T periods	Section 2
$b_T(s)$	premium when seller's score is s	Section 2
$R(q, T, \delta)$	minimum ratio (over all rating histories and all possible deviations) of future decrease in payoff to current gain in payoff, if seller deviates from being truthful in the current period	Section 3
$T^*(\delta)$	optimal window size given δ	Section 3
$V(\vec{r})$	seller's maximal expected discounted profit after rating history \vec{r}	Sections 4, 5, 6
$T_p^*(\delta)$	optimal window size given δ under perfect monitoring	Section 4
$T_m^*(\delta)$	optimal window size given δ with missing ratings	Section 5
$T_w^*(\delta)$	optimal window size given δ with inaccurate ratings	Section 6

Table 1 Notation used in the paper.

(Section 4), missing ratings (Section 5), and inaccurate ratings (Section 6). In each of these settings our approach is to show that the two properties required in Theorem 1 hold, and thus the optimal window is increasing in δ and decreasing in q .

3.1. Optimization

We assume that the goal of the mechanism designer is to maximize the range of parameters for which it is optimal for the seller to be always truthful. The optimal window choice depends on the available information regarding δ and q . In particular, if the mechanism designer knows δ but not q , a reasonable choice of T is one which maximizes the range of values of q for which the seller will be always truthful. On the other hand, if the mechanism designer knows q but not δ , a reasonable choice of T is one which maximizes the range of values of δ for which the seller will be always truthful.

We provide sufficient conditions (Properties 1 and 2 below) to ensure that (1) the window size

that maximizes the range of q for which the seller is always truthful is *increasing in δ* , and (2) the window size that maximizes the range of δ for which the seller is always truthful is *decreasing in q* . In the following sections we verify these sufficient conditions for perfect monitoring (Section 4), imperfect monitoring with missing ratings (Sections 5), and imperfect monitoring with inaccurate ratings (Section 6). Thus, we conclude that the optimal window is increasing in δ and decreasing in q for a wide range of settings.

Our analysis proceeds as follows. If the seller deviates from being truthful in the current period, then his current payoff increases while his future payoff decreases. The seller will not deviate from being truthful if the future decrease exceeds the current gain. Let $R(q, T, \delta)$ be the *minimum possible ratio* (over all rating histories and all possible deviations) of future decrease over current gain, if the seller deviates from being truthful in the current period. Note that by definition of R , *it is optimal for the seller to be always truthful if and only if $R(q, T, \delta) \geq 1$* .

We now introduce the two key properties that are critical to our analysis.

Property 1 *There exists a function $F(T, \delta)$ that is increasing in δ for each fixed T , such that*

$$R(q, T, \delta) = q \cdot F(T, \delta).$$

Note that if Property 1 holds, then it is optimal for the seller to always be truthful if and only if $q \cdot F(T, \delta) \geq 1$. If Property 1 holds, define the set

$$T^*(\delta) = \arg \max_T F(T, \delta)$$

for each δ . For tractability, we assume that for all δ , the set $T^*(\delta)$ is nonempty.

We now introduce the second property.

Property 2 *The set $T^*(\delta)$ is increasing in δ in the following sense: for $\delta \geq \delta'$,*

$$(i) \max\{T : T \in T^*(\delta)\} \geq \max\{T : T \in T^*(\delta')\}; \text{ and}$$

$$(ii) \min\{T : T \in T^*(\delta)\} \geq \min\{T : T \in T^*(\delta')\}.$$

We next show the main result of this section: Properties 1 and 2 are sufficient conditions to ensure that the optimal window is increasing in δ and decreasing in q . This theorem provides

a general tool we can employ to analyze specific settings: the task of characterizing behavior of optimal windows has been reduced to verifying that Properties 1 and 2 hold.

THEOREM 1. *Suppose that Properties 1 and 2 hold.*

(i) *If δ is known and the goal is to maximize the range of q for which the seller is always truthful, then the set of optimal windows is increasing in δ .*

(ii) *If q is known and the goal is to maximize the range of δ for which the seller is always truthful, then the set of optimal windows is decreasing in q .*

Proof: (i) *The mechanism designer knows δ , but not q .* The goal is to maximize the range of q for which the seller is always truthful. From Property 1, it is optimal for the seller to be always truthful if and only if $q \cdot F(T, \delta) \geq 1$. This implies we should maximize $F(T, \delta)$, i.e., any $T \in T^*(\delta)$ is an optimal choice. We conclude that if $T^*(\delta)$ is increasing (Property 2), then *the set of optimal windows is increasing in δ .*

(ii) *The mechanism designer knows q , but not δ .* The goal is to maximize the range of δ for which the seller is always truthful. From Property 1, it is optimal for the seller to be always truthful if and only if $q \cdot F(T, \delta) \geq 1$. Thus, given q , we solve:

$$\begin{aligned} & \text{minimize } \delta \\ & \text{subject to } q \cdot F(T, \delta) \geq 1 \end{aligned}$$

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* window size. It then follows that any $T \in T^*(\delta^*(q))$ is an optimal choice. Observe that since the constraint is increasing in q and δ , it follows that $\delta^*(q)$ is decreasing in q . We conclude that if $T^*(\delta)$ is increasing (Property 2), then *the set of optimal windows is decreasing in q .* In words, as q increases, it is possible to make sellers with smaller δ truthful, and for such sellers a smaller window size is more appropriate. ■

We can now see the insight discussed at the beginning of the section. First, increasing the window size is more likely to make patient sellers (those with high δ) truthful. This is an intuitive result,

since sellers with larger δ are more patient, and thus an aggregation mechanism with longer memory can successfully couple current behavior with distant future payoffs. On the other hand, a larger window 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.

By choosing a window size T , the system designer determines *how much* and for *how long* the seller's future scores decrease if he does not describe his current item truthfully. In particular, if in the current period the seller receives a rating that is smaller than the maximum possible rating by d , then his score will decrease by d/T in each of the next T periods (relative to receiving the maximum possible rating). We have the following tradeoff (noted in the Introduction) between the intensity and the duration of this score reduction: the intensity is decreasing in T , while the duration is increasing in T . The optimal value of T will depend on the available information on the seller's attributes and the premium function. As δ increases the duration effect becomes more important, so the optimal window size increases. On the other hand, as q increases the intensity effect becomes more important, so the optimal window size decreases. Crucially, we observe that this tradeoff is also faced by a mechanism designer who knows *neither* δ nor q : a choice must be made regarding the incentives provided to patient sellers and those provided to high quality sellers.

This paper focuses on the set of *optimal* window sizes as defined in Theorem 1. In words, this is the set of windows that maximize the range of seller parameters for which it is optimal for the seller to be always truthful. We conclude this section by noting that the set of *feasible* window sizes under which the seller is truthful has different properties than the set of *optimal* windows. In particular, under Property 1, the seller is always truthful under any window in the set $\{T : q \cdot F(T, \delta) \geq 1\}$. The following lemma characterizes properties of this set. The intuition behind this lemma is that it is easier to incentivize a more patient seller.

LEMMA 1. *Suppose that Property 1 holds and given q , let $\tilde{T}(\delta) \equiv \{T : q \cdot F(T, \delta) \geq 1\}$. Then $\min \tilde{T}(\delta)$ is decreasing in δ , and $\max \tilde{T}(\delta)$ is increasing in δ .*

3.2. Roadmap

Theorem 1 demonstrates a tradeoff between incentivizing patient and high quality sellers: if Properties 1 and 2 hold, then the optimal window is increasing in δ and decreasing in q . In the following sections we show this tradeoff for a range of settings. First, we consider the case of perfect monitoring (Section 4), where the buyer rates the seller accurately after every transaction. Then, we consider two types of imperfect monitoring: the seller may not receive ratings after some transactions (Section 5), and the rating that the seller receives may not accurately reflect his action (Section 6).

As suggested by Theorem 1, our analysis in each of the three subsequent sections consists of the following steps.

1. First, we verify Property 1: we show there exists a function $F(T, \delta)$ that is increasing in δ such that $R(q, T, \delta) = q \cdot F(T, \delta)$, so that it is optimal for the seller to be always truthful if and only if $q \cdot F(T, \delta) \geq 1$. A key requirement for this step is that the premium function is logarithmically concave (Assumption 2).

2. Second, we verify Property 2: we show that the set $T^*(\delta) = \arg \max_T \{F(T, \delta)\}$ is increasing in δ .

3. Finally, we apply Theorem 1 to conclude that the optimal window is increasing in δ and decreasing in q .

4. Perfect Monitoring

In this section we apply the framework of Section 3 to the case of perfect monitoring. In particular, we assume that the seller receives a rating that accurately reflects his action after every transaction.

We assume that the rating that the seller receives depends on the difference between the description that the seller posted (v_d), and the true value of the item (v). In particular, we assume that the rating is equal to $\alpha - \beta(v_d - v)^+$, where $x^+ \equiv \max\{x, 0\}$ is the positive part of x , and $0 \leq \beta(v_H - v_L) \leq \alpha$. (Recall that $0 \leq v_L < v_H \leq 1$.) That is, the seller receives a low rating if he exaggerated the value of the item in his description, but not if he understated it. The seller receives the best possible rating (i.e., α) if he did not exaggerate the value of the item in his description. This is a reasonable assumption; however, it is not essential: for instance, our results also hold

if the rating is equal to $\alpha - \beta|v_d - v|$. For the remainder of the section, we assume without loss of generality that $\alpha = \beta = 1$; the results continue to hold for any values of α and β (as long as $0 \leq \beta(v_H - v_L) \leq \alpha$).

For this section we assume that the value of the item X takes values in $[v_L, v_H]$ according to a distribution with continuous density f with support $[v_L, v_H]$. The extension to discrete probability distributions is straightforward.

Let $V(\vec{r})$ be the maximum infinite horizon discounted payoff of the seller when his current vector of ratings is \vec{r} . Since the expected payment to the seller is increasing in his description, it is never optimal for the seller to understate the value of an item.

The seller's optimal policy is given by the following dynamic program.

$$V(\vec{r}) = \int_{v_L}^{v_H} \max_{y \in [x, v_H]} \left\{ y \cdot b_T \left(\sum_{i=1}^T r_i / T \right) + \delta \cdot V(1 - (y - x)^+, \vec{r}) \right\} f(x) dx \quad (1)$$

In particular, if the value of the item is x , the seller chooses a description that maximizes his infinite horizon discounted payoff. If the seller describes the item truthfully, his payoff is $x \cdot b_T(\sum_{i=1}^T r_i / T) + \delta V(1, \vec{r})$, since he receives $x \cdot b_T(\sum_{i=1}^T r_i / T)$ and his ratings "increase" to $(1, \vec{r})$. If he describes it as an item of value $y > x$, his payoff is $y \cdot b_T(\sum_{i=1}^T r_i / T) + \delta V(1 - (y - x), \vec{r})$, since he receives $y \cdot b_T(\sum_{i=1}^T r_i / T)$ now, but his ratings "decrease" to $(1 - (y - x), \vec{r})$. The seller will choose the description with the maximum payoff.

We say that the seller is *truthful at \vec{r}* if it is optimal for him to describe an item of value x truthfully when his rating vector is \vec{r} for all $x \in [v_L, v_H]$ with $f(x) > 0$. By (1), it is optimal for the seller to be truthful at \vec{r} if and only if

$$(y - x) \cdot b_T \left(\sum_{i=1}^T r_i / T \right) \leq \delta (V(1, \vec{r}) - V(1 - (y - x), \vec{r})) \quad (2)$$

for all $y \in [x, v_H]$. In particular, if the seller is untruthful and posts a description $y > x$, then his current payoff will increase by $(y - x) \cdot b_T \left(\sum_{i=1}^T r_i / T \right)$ but his expected payoff starting from the next period will decrease by $V(1, \vec{r}) - V(1 - (y - x), \vec{r})$ (relative to being truthful).

PROPOSITION 1. *Property 1 holds under perfect monitoring: in particular, there holds $R(q, T, \delta) = q \cdot F_p(T, \delta)$, where*

$$F_p(T, \delta) = \left(\min_{d \in (0, v_H - v_L]} \left\{ \frac{b_T(1) - b_T(1 - d/T)}{d \cdot b_T(1)} \right\} \right) \sum_{i=1}^T \delta^i.$$

Thus it is optimal for the seller to be truthful at all \vec{r} if and only if q , δ , and T jointly satisfy the following constraint:

$$q \cdot F_p(T, \delta) \geq 1.$$

We use the subscript p to denote that we are considering perfect monitoring, i.e., that the seller receives a rating that accurately reflects his action after every transaction.

Proposition 1 reduces the problem of checking whether it is optimal for the seller to be always truthful to checking whether the following inequality holds for all $d \in [0, v_H - v_L]$.

$$d \cdot b_T(1) \leq (b_T(1) - b_T(1 - d/T)) \sum_{i=1}^T \delta^i$$

This condition ensures that the seller does not deviate from being truthful when his score is $s = 1$. This is the case if expected future gains for being truthful are greater than current gains for being untruthful. In particular, if the seller exaggerates his description by d , then his current payment increases by $d \cdot b_T(1)$, but future payments decrease by $b_T(1) - b_T(1 - d/T)$ in each of the T next periods.

Our analysis depends on analyzing the set $T_p^*(\delta)$ defined as follows for each δ :

$$T_p^*(\delta) = \arg \max_{T \geq 1} F_p(T, \delta).$$

We next show that Property 2 holds under perfect monitoring.

PROPOSITION 2. *Property 2 holds under perfect monitoring, i.e., $T_p^*(\delta)$ is increasing in δ .*

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

Using Propositions 1 and 2, we can apply Theorem 1 in the case of perfect monitoring:

(1) If δ is known and the goal is to maximize the range of q for which the seller is always truthful, then the set of optimal windows is increasing in δ ; and

(2) If q is known and the goal is to maximize the range of δ for which the seller is always truthful, then the set of optimal windows is decreasing in q .

5. Missing Ratings

The previous section assumed perfect monitoring, i.e., that the seller receives a rating which accurately reflects his action after every transaction. However, various studies have shown that monitoring may be imperfect in practice (e.g., Dellarocas and Wood 2008, Chwelos and Dhar 2008, Bolton et al. 2009). In this section and the following section we relax the assumption on perfect monitoring in two ways. First, in this section, we consider the case of missing ratings, where after some transactions the buyer does not rate the seller. In the next section, we consider the case where ratings may not always reflect the seller's action. In this section we show that in the presence of missing ratings, the optimal window size is increasing in δ and decreasing in q .

For simplicity, we focus on a binary setting: we assume that in every period the seller has an item for sale which has value v_H with probability q_H and value v_L with probability $1 - q_H$, where $0 \leq v_L < v_H$. As before, let q be the expected value of the item that the seller has for sale; in this setting $q = q_H v_H + (1 - q_H) v_L$. In the beginning of every period, the seller observes the item he has for sale and chooses a description of v_L or v_H .

Let p_{av} be the probability that the seller receives no rating when his action is $a \in \{t, u\}$ for being truthful and untruthful respectively, and the true value of the item is $v \in \{v_H, v_L\}$. If the seller receives a rating, we assume that it accurately reflects his action: he receives a good rating (of value 1) for describing his item truthfully, and a bad rating (of value 0) for exaggerating the value of a low value item in his description. We assume that p_{tH} is not significantly larger than p_{uH} , so that it is optimal for the seller to describe a high value item truthfully.

Let $V(\vec{r})$ be the maximum infinite horizon discounted value when the current vector of ratings is \vec{r} ; then:

$$\begin{aligned}
 V(\vec{r}) = & q_H \left(v_H \cdot b_T \left(\sum_{i=1}^T r_i / T \right) + \delta((1 - p_{tH})V(1, \vec{r}) + p_{tH}V(\vec{r})) \right) \\
 & + (1 - q_H) \max \left\{ v_H \cdot b_T \left(\sum_{i=1}^T r_i / T \right) + \delta((1 - p_{uL})V(0, \vec{r}) + p_{uL}V(\vec{r})), \right. \\
 & \left. v_L \cdot b_T \left(\sum_{i=1}^T r_i / T \right) + \delta((1 - p_{tL})V(1, \vec{r}) + p_{tL}V(\vec{r})) \right\} \quad (3)
 \end{aligned}$$

In particular, with probability q_H the seller has a high value item for sale, which he describes truthfully. The immediate payment he receives is $v_H \cdot b_T(\sum_{i=1}^T r_i / T)$; with probability $1 - p_{tH}$ he

receives a good rating and with probability p_{tH} he receives no rating. With probability $1 - q_H$ the seller has a low value item for sale. If he describes it as a high value item, his payoff is $v_H \cdot b_T(\sum_{i=1}^T r_i/T) + \delta((1 - p_{uL})V(0, \vec{r}) + p_{uL}V(\vec{r}))$, since he receives $v_H \cdot b_T(\sum_{i=1}^T r_i/T)$ now, but his ratings “decrease” to $(0, \vec{r})$ with probability $1 - p_{uL}$ and remain the same with probability p_{uL} . If he describes the item truthfully, he receives a lower payment now, but his ratings “increase” to $(1, \vec{r})$ with probability $1 - p_{tL}$. The seller will choose the description that maximizes his payoff.

As in Section 3, our goal is to maximize the range of parameters for which is it optimal for the seller to be truthful. Let $p \equiv q_H \cdot p_{tH} + (1 - q_H) \cdot p_{tL}$ be the ex ante probability (before the value of the item is known) that the seller receives no rating if he is truthful.

PROPOSITION 3. *If $p_{uL} \geq p_{tL}$, then Property 1 holds with missing ratings: in particular, $R(q, T, \delta) = q \cdot F_m(T, \delta)$, where*

$$F_m(T, \delta) = \frac{1 - p_{uL}}{v_H - v_L} \frac{b_T(1) - b_T(1 - 1/T)}{b_T(1)} \sum_{i=0}^{\infty} \delta^{i+1} \sum_{j=0}^{\min(T-1, i)} \binom{i}{j} (1-p)^j p^{i-j}.$$

Thus it is optimal for the seller to be truthful at all \vec{r} if and only if q , δ , and T jointly satisfy the following constraint:

$$q \cdot F_m(T, \delta) \geq 1.$$

We use the subscript m to denote that we are considering the possibility of missing ratings.

Proposition 3 reduces the problem of finding whether it is optimal for the seller to be always truthful to checking whether the following inequality holds.

$$(v_H - v_L)b_T(1) \leq (1 - p_{uL})q(b_T(1) - b_T(1 - 1/T)) \sum_{i=0}^{\infty} \delta^{i+1} \sum_{j=0}^{\min(T-1, i)} \binom{i}{j} (1-p)^j p^{i-j}$$

This condition requires that the seller does not deviate from being truthful when his current score is equal to 1. The seller does not deviate from being truthful if the discounted future gains for being truthful are greater (in expectation) than the current gains for being untruthful. The current payment to the seller increases by $(v_H - v_L)b_T(1)$ if the seller deviates from being truthful. On the other hand, the future payments to the seller decrease by $b_T(1) - b_T(1 - 1/T)$ in every period that is affected by the current rating. This is the case until the seller receives T new ratings. The seller

receives exactly j new ratings in i periods with probability $(1-p)^j p^{i-j}$. Therefore, the probability that the seller has not received T new ratings in i periods is $\sum_{j=0}^{\min(T-1,i)} \binom{i}{j} (1-p)^j p^{i-j}$.

Let $T_m^*(\delta) = \arg \max_T \{F_m(T, \delta)\}$.

PROPOSITION 4. *Property 2 holds with missing ratings, i.e., $T_m^*(\delta)$ is increasing in δ .*

Propositions 3 and 4 give conditions under which Properties 1 and 2 hold under missing ratings.

We apply Theorem 1 to conclude in the case of missing ratings, if p is fixed and $p_{uL} \geq p_{tL}$:

(1) If δ is known and the goal is to maximize the range of q for which the seller is always truthful, then the set of optimal windows is increasing in δ ; and

(2) If q is known and the goal is to maximize the range of δ for which the seller is always truthful, then the set of optimal windows is decreasing in q .

We note that we are assuming that p remains fixed as δ and q change. Since $p = q_H \cdot p_{tH} + (1 - q_H) \cdot p_{tL}$ and $q = q_H \cdot v_H + (1 - q_H) \cdot v_L$, we can assume that q_H is fixed and q changes through v_H and v_L . Alternatively, we can assume that q_H is changing, but also p_{tH} and p_{tL} change accordingly so that p is fixed.

We conclude by discussing two cases that can be viewed as special cases of (3). First, consider the setting where the item that the seller has for sale is not always sold. If the probability that the item is not sold depends on the description that the seller posts, but not on his score, then we can use (3) and still interpret $b_T(s)$ as the expected premium to the seller.

Second, we note that our current model aggregates ratings from rounds in which the seller has a low value item in exactly the same way as ratings from rounds in which the seller has a high value item. This modeling decision was motivated by the way ratings are aggregated in online markets. However, the seller faces no moral hazard for describing a high value item, and it may be reasonable to assume that the seller is not given a rating for advertising a high value item truthfully. Then we can assume that $p_{tH} = 1$ in (3) and still conclude that the optimal window is increasing in δ and decreasing in q .

6. Inaccurate Ratings

The rating that the seller receives may not reflect his action, because the buyer may make a mistake. In this section we identify conditions under which in the presence of inaccurate ratings

the optimal window is increasing in δ and decreasing in q .

As in Section 5 we consider a binary setting. By a slight abuse of notation, we model inaccurate ratings by assuming that with probability p_{av} the seller receives the wrong rating when his action is $a \in \{t, u\}$ for being truthful and untruthful respectively, and the true value of the item is $v \in \{v_H, v_L\}$. We assume that the probability p_{tH} is sufficiently small so that the seller always describes a high value item truthfully. Let $V(\vec{r})$ be the maximum infinite horizon discounted payoff when the current vector of ratings is \vec{r} . Then, the optimization problem of the seller is given by the following dynamic program.

$$\begin{aligned} V(\vec{r}) = & q_H \left(v_H \cdot b_T \left(\sum_{i=1}^T r_i/T \right) + \delta \left((1 - p_{tH})V(1, \vec{r}) + p_{tH}V(0, \vec{r}) \right) \right) \\ & + (1 - q_H) \max \left\{ v_H \cdot b_T \left(\sum_{i=1}^T r_i/T \right) + \delta \left((1 - p_{uL})V(0, \vec{r}) + p_{uL}V(1, \vec{r}) \right), \right. \\ & \left. v_L \cdot b_T \left(\sum_{i=1}^T r_i/T \right) + \delta \left((1 - p_{tL})V(1, \vec{r}) + p_{tL}V(0, \vec{r}) \right) \right\} \end{aligned}$$

In particular, if the seller describes a low value item as a high value item, he receives a bad rating with probability $1 - p_{uL}$ and a good rating with probability p_{uL} . On the other hand, if the seller has a low value item for sale and describes it truthfully, then he receives an immediate payment of $v_L \cdot b_T \left(\sum_{i=1}^T r_i/T \right)$ in expectation, and gets a good rating with probability $1 - p_{tL}$; with probability p_{tL} he gets a bad rating despite the fact that he described the item truthfully.

Let $p \equiv q_H \cdot p_{tH} + (1 - q_H) \cdot p_{tL}$ be the ex ante probability (before the value of the item is known) that the seller receives a negative rating if he is truthful.

PROPOSITION 5. *If p is sufficiently small, then Property 1 holds with inaccurate ratings: in particular, $R(q, T, \delta) = q \cdot F_w(T, \delta)$, where*

$$F_w(T, \delta) = \frac{1 - p_{uL} - p_{tL}}{v_H - v_L} \sum_{i=0}^{T-1} \delta^{i+1} \sum_{k=0}^i \binom{i}{k} p^k (1-p)^{i-k} \frac{b_T(1 - k/T) - b_T(1 - (k+1)/T)}{b_T(1)}.$$

Thus it is optimal for the seller to be truthful at all \vec{r} if and only if q , δ , and T jointly satisfy the following constraint:

$$q \cdot F_w(T, \delta) \geq 1.$$

We use the subscript w to denote that we are considering the possibility of wrong ratings, i.e., that the rating may not accurately reflect the seller's action.

Proposition 5 reduces the problem of finding whether it is optimal for the seller to be always truthful to checking whether the following inequality holds.

$$(v_H - v_L)b_T(1) \leq (1 - p_{uL} - p_{tL})q \sum_{i=0}^{T-1} \delta^{i+1} \sum_{k=0}^i \binom{i}{k} p^k (1-p)^{i-k} \frac{b_T(1 - k/T) - b_T(1 - (k+1)/T)}{b_T(1)}$$

This inequality checks whether the seller would deviate from being truthful if his current score is equal to 1 and he has a low value item for sale. The seller does not deviate from being truthful if the discounted future gains for being truthful are greater (in expectation) than the current gains for being untruthful. If the seller deviates from being truthful by describing a low value item as a high value item, then his current payment increases by $(v_H - v_L)b_T(1)$. If the seller gets a good rating now and is truthful in future periods, then in i periods from now his score is $1 - k/T$ with probability $\binom{i}{k} p^k (1-p)^{i-k}$. On the other hand, if the seller gets a bad rating now and is truthful in future periods, then in i periods from now his score is $1 - (k+1)/T$ with probability $\binom{i}{k} p^k (1-p)^{i-k}$. The right hand side considers the difference in the expected payments in the next T periods, and discounts appropriately.

We conjecture that the condition that p must be sufficiently small in Proposition 5 can be weakened.³ Numerical experiments with specific premium functions suggest that the result of Proposition 5 holds for any value of p . However, technical verification of this fact in general remains an open problem.

According to Theorem 1, in order to conclude that the optimal window is increasing in δ and decreasing in q , we verify Property 2: we derive conditions under which the set

$$T_w^*(\delta) = \arg \max_T \{F_w(T, \delta)\}$$

is increasing in δ . This is done in the following proposition.

³ The proof of Proposition 5 shows that Property 1 holds with inaccurate ratings for arbitrary values of p when b_T is either concave or exponential.

PROPOSITION 6. *Let*

$$h_{T,T'}(k) = \frac{b_{T'}(1 - k/T')}{b_{T'}(1)} - \frac{b_T(1 - k/T)}{b_T(1)}.$$

Suppose that for every $T' > T$ there exists an integer k_0 , $0 \leq k_0 \leq T$, such that $h_{T,T'}(k) \leq h_{T,T'}(k+1)$ for $k < k_0$ and $h_{T,T'}(k) \geq h_{T,T'}(k+1)$ for $k > k_0$.

Then Property 2 holds with inaccurate ratings, i.e., $T_w^(\delta)$ is increasing in δ .*

Propositions 5 and 6 provide conditions under which Properties 1 and 2 hold. We can use Theorem 1 to conclude that if the condition of Proposition 6 holds, then:

(1) If δ is known and the goal is to maximize the range of q for which the seller is always truthful, then for any $\bar{\delta} < 1$ and for sufficiently small p , the optimal window is increasing in δ in the interval $[0, \bar{\delta}]^4$; and

(2) If q is known and the goal is to maximize the range of δ for which the seller is always truthful, for every range of q and for sufficiently small p the optimal window is decreasing in q .

The condition of Proposition 6 is satisfied by many premium functions. As an example, consider premium functions of the form $b_T(s) = \alpha(T) \cdot b(s) + \gamma(T)$, where $\alpha(T)$ is nondecreasing in T and $\gamma(T)$ is nonincreasing in T . This form captures the following intuition: *as the window size increases, buyers trust the information that is aggregated in the seller's score more*. For instance, if the seller has the maximum possible score ($s = 1$), we expect that the premium will increase as T increases; this is captured by the assumption that $\alpha(T)$ is increasing. On the other hand, if the seller has the minimum possible score ($s = 0$), we expect the premium to decrease as T increases; this is captured by the assumption that $\gamma(T)$ is nonincreasing. Simple calculations show that the condition of Proposition 6 is satisfied for various functions of this form; e.g., if $b(s) = s^n$ or $b(s) = e^{ns}$, and $\alpha(\cdot)$, $\gamma(\cdot)$ are arbitrary functions of T . We note that many empirical studies use regression forms that correspond to premia of the form $\alpha(T) \cdot b(s) + \gamma(T)$, where $b(s) = e^{ns}$ (e.g., Cabral and Hortacsu 2009, Lucking-Reiley et al. 2007, Ghose et al. 2005).

The following corollary restricts attention to premia that do not explicitly depend on the window size.

⁴The upper bound on p is some increasing function of T , say $u(T)$. Consider $\bar{\delta} < 1$. If b_T satisfies the condition of Proposition 6, then the optimal window is increasing for $\delta \in [0, \bar{\delta}]$ if $p \leq u(\max\{T_w^*(\bar{\delta})\})$.

COROLLARY 1. Suppose $b_T(\cdot) \equiv b(\cdot)$. If $b'(s)$ is logarithmically concave, then Property 2 holds with inaccurate ratings, i.e., $T_w^*(\delta)$ is increasing in δ .

Examples of functions with a logarithmically concave derivative are $b(s) = s^n$, $b(s) = e^s$ and the logistic function $b(s) = 1/(1 + e^{a-bs})$ for $b > 0$. We note that the conclusion of Corollary 1 holds more generally if $b'(1-y) - yb''(1-y) < 0$ for some $y \in [0, 1]$ implies that $b'(1-z) - zb''(1-z) < 0$ for $z > y$.

We conclude by summarizing the results of this section. In the case of inaccurate ratings, if p is small and the condition of Proposition 6 holds, then assuming that p is fixed: (1) if δ is known and the goal is to maximize the range of q for which the seller is always truthful, then the set of optimal windows is increasing in δ ; (2) if q is known and the goal is to maximize the range of δ for which the seller is always truthful, then the set of optimal windows is decreasing in q .

7. Windows Measured in Real Time

This paper considers the dependence of the optimal window on the seller's discount factor δ and the average value of the items for sale q . We have shown that the optimal window is increasing in δ and decreasing in q for a range of settings. In this section we briefly discuss the dependence of the optimal window on the rate at which the seller transacts and the steepness of the premium function.

Our analysis throughout the paper considers the optimal number of transactions that should be included in the seller's score. However, electronic marketplaces usually provide information for some fixed window of *real time*. As long as sellers transact at a fixed rate per unit time, the optimal window is increasing in δ and decreasing in q even when the goal is to choose the optimal window measured in real time (e.g., months or days).

We now investigate the dependence of the optimal window on the rate at which a seller transacts, and on the steepness of the premium function. We assume that the expected premium to the seller as a function of his score is $b(s) = s^n$, so that n characterizes the steepness of the premium function; we have observed similar qualitative behavior with other choices of premia. Moreover, we assume that all sellers have the same annual discount factor δ . Sellers may sell items at different




Figure 1 Contour plots of the number of months that maximizes the range of q for which the seller is always truthful. We assume that the premium function is $b(s) = s^n$. Each plot considers a different annual discount factor.

frequencies. A seller that is selling f items per year discounts the payment he will receive from the next item that he will have for sale by $\delta^{1/f}$. Assuming that the annual discount factor is known, we wish to find the number of months that maximizes the range of q for which the seller is always truthful. (The same qualitative insights hold if q is fixed and the goal is to maximize the range of δ for which the seller is always truthful.)

Figure 1 shows contour plots of the optimal number of months for three values of the annual discount factor; recall that twelve months is the window size chosen by both eBay and the Amazon Marketplace. The horizontal axis represents the steepness of the premium function, and the vertical axis shows how many items the seller sells per year. The plots demonstrate that the optimal number of months increases as the discount factor increases (for a fixed point). Moreover, Figure 1 suggests that *the optimal number of months increases as the premium function becomes steeper, and as the number of items per year decreases*. Both these observations match intuition. Clearly, if a seller sells fewer items per year, then the window (in real time) should be larger to ensure a sufficiently large number of transactions is considered. Further, if the premium function is very sensitive to the seller's reputation, then a larger window can be used while still providing incentives for the seller to be truthful.

8. Conclusions

This paper studies the design of Window Aggregation Mechanisms, i.e., mechanisms that show the average value of the seller's ratings within some fixed window of past transactions. With the goal of incentivizing the seller to describe the item truthfully, we address the design question of choosing the right window size. The main result is an interesting qualitative tradeoff: increasing the window

size is more likely to make patient sellers truthful, while it is less likely to make high quality sellers truthful. This result provides suggestions for designing rating aggregation mechanisms.

First, the optimal window size is increasing in the discount factor of the seller. Since interest rates are decreasing in discount factors, we conclude that the optimal window size is decreasing in the interest rate. That is, as interest rates decrease, larger window sizes become more appropriate. This intuition holds regardless of whether the window size measures the number of transactions or real time.

Second, the optimal window size is decreasing in the seller's quality. In particular, if the seller's quality is known and we wish to maximize the range of discount factors for which it is optimal for the seller to be truthful, then the optimal window size is decreasing in the seller's quality. This suggests that larger window sizes are more appropriate for marketplaces where sellers have mostly low value items.

Almost all prominent electronic marketplaces, such as eBay and Amazon, use variants of the Window Aggregation Mechanism. Our main insight suggests these marketplaces face a tradeoff between incentivizing patient sellers and incentivizing high quality sellers: larger windows are more likely to make the former truthful, but less likely to make the latter truthful. However, this insight is predicated on the assumption that the marketplace does not distinguish between sub-markets for items of differing values.

Given the heterogeneity of sellers and items in marketplaces such as eBay and Amazon, the aforementioned discussion suggests that Window Aggregation Mechanisms could be adapted to the target application. This raises the following intriguing question: is it better to use different window sizes for different markets (i.e., markets for items of different value) within the same marketplace? The answer depends on the percentage of sellers that sell across markets. Alternatively, it could be beneficial to adjust the window size over the lifetime of the seller (e.g., depending on the value of items he is selling) or to discount past ratings in some more general way than the Window Aggregation Mechanism. Such mechanisms have been considered in the literature (e.g., Fan et al. 2005, Aperjis and Johari 2008); however, finding the best aggregation mechanism in this broader sense remains an open question.

Acknowledgments. This work was partially supported by the National Science Foundation under grant 0428868. We are grateful to the Editor, Associate Editor and two anonymous referees for extremely helpful comments. We also gratefully acknowledge helpful conversations with Ashish Goel, Robert Kleinberg, Richard Zeckhauser, Benjamin Van Roy, Nimrod Megiddo, Sunil Kumar, and Nick Bambos.

References

- Akerlof, George. 1970. The market for lemons: Quality uncertainty and the market mechanism. *Quart. J. of Econom.* **84**(3) 488–500.
- Aperjis, Christina, Ramesh Johari. 2008. Reputation mechanisms and efficient trade. Tech. rep., Stanford University.
- Bolton, Gary, Ben Greiner, Axel Ockenfels. 2009. Engineering trust - reciprocity in the production of reputation information. Working paper series in economics, University of Cologne, Department of Economics.
- Cabral, Luis, Ali Hortacsu. 2009. Dynamics of seller reputation: Theory and evidence from eBay. *J. of Industr. Econom.* (to appear) .
- Chwelos, Paul, Tirtha Dhar. 2008. Differences in “truthiness” across online reputation mechanisms. Working Paper, Sauder School of Business.
- Cripps, Martin, George 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, Charles Wood. 2008. The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias. *Management Sci.* **54**(3) 460–476.
- Fan, Ming, Yong Tan, Andrew 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 Ipeirotis, Arun Sundararajan. 2005. Reputation premiums in electronic peer-to-peer markets: Analyzing textual feedback and network structure. *P2PECON*. 150–154.

- Li, Lingfang. 2009. Reputation, trust, and rebates: How online auction markets can improve their feedback systems. *J. of Econom. and Manag. Strat.*, (to appear) .
- 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.
- Milgrom, Paul, Chris Shannon. 1994. Monotone comparative statics. *Econometrica* **62**(1).
- 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. 1998. *Supermodularity and complementarity*. Princeton University Press.
- Tversky, Amos, Daniel Kahneman. 1974. Judgment under uncertainty: Heuristics and biases. *Science* **185** 1124–1131.

Appendix: Proofs

This appendix contains the proofs of the lemmas and propositions of the paper. We first give a result that is used in multiple proofs.

LEMMA 2. *Let $f : [0, 1] \rightarrow R_+$ be an increasing and logarithmically concave function. Then*

$$(f(1) - f(1 - t)) \cdot f(k) \leq f(1) \cdot (f(k) - f(k - t)) \quad (4)$$

for all $k < 1$ and $t \in (0, k]$.

Proof of Lemma 2: Condition (4) is equivalent to

$$f(1 - t) \cdot f(k) \geq f(1) \cdot f(k - t).$$

Since f is nonnegative, this trivially holds if $f(k - t) = 0$. Moreover, if either $f(1 - t) = 0$ or $f(k) = 0$, then $f(k - t) = 0$ (since f is increasing).

If $f(k - t) > 0$, it suffices to show that $f(x)/f(x - t)$ is nonincreasing in x for $x \in [k, 1]$. Since f is logarithmically concave, we conclude that $\log(f(x)) - \log(f(x - t))$ is nonincreasing in x , which implies that $f(x)/f(x - t)$ is nonincreasing in x . This concludes the proof. ■

Proof of Lemma 1: Let $\delta_1 < \delta_2$. If $T_1 \in \tilde{T}(\delta_1)$, then $q \cdot F(T_1, \delta_1) \geq 1$. Since $F(T, \delta)$ is increasing in δ (by Property 1), we conclude that $q \cdot F(T_1, \delta_2) \geq 1$ and thus $T_1 \in \tilde{T}(\delta_2)$. This implies that $\tilde{T}(\delta_1) \subseteq$

$\tilde{T}(\delta_2)$ whenever $\delta_1 < \delta_2$. It follows that $\min \tilde{T}(\delta)$ is decreasing in δ , and $\max \tilde{T}(\delta)$ is increasing in δ . ■

Proof of Proposition 1: $R(q, T, \delta)$ is the minimum ratio (over all rating histories and all possible deviations) of future decrease in payoff to current gain in payoff, if the seller deviates from being truthful in the current period. In other words, it is optimal for the seller to be always truthful if and only if $R(q, T, \delta) \geq 1$ (by the definition of $R(q, T, \delta)$). Thus, to show that $R(q, T, \delta) = q \cdot F_p(T, \delta)$ it suffices to show that it is optimal for the seller to be always truthful if and only if $q \cdot F_p(T, \delta) \geq 1$.

We proceed by showing the latter in this proof.

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. The seller may consider to exaggerate an item in his description when the value of the item is $v < v_H$. 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 he gets a rating of value 1 in all future periods, given that his current ratings are \vec{r} . Note that $s_0(\vec{r}) = \sum_{i=1}^T r_i/T$ is the seller's current score. Then,

$$\hat{V}(1, \vec{r}) - \hat{V}(\rho, \vec{r}) = q \sum_{i=0}^{T-1} \delta^i (b_T(s_i(1, \vec{r})) - b_T(s_i(\rho, \vec{r})))$$

and $s_i(1, \vec{r}) - s_i(\rho, \vec{r}) = (1 - \rho)/T$. By (2), the seller is not better off deviating from the truthful policy when his ratings are \vec{r} if and only if

$$d \cdot b_T(s_0(\vec{r})) \leq \delta (\hat{V}(1, \vec{r}) - \hat{V}(1 - d, \vec{r}))$$

for $d \in [0, D]$, where $D = v_H - v_L$. It is optimal for the seller to be always truthful if and only if the previous condition holds for all \vec{r} . Substituting $\hat{V}(1, \vec{r}) - \hat{V}(\rho, \vec{r})$ in the previous condition and setting $d = 1 - \rho$, we conclude that it is optimal for the seller to be truthful at all \vec{r} if and only if

$$d \leq q \sum_{i=0}^{T-1} \delta^{i+1} \frac{b_T(s_i(1, \vec{r})) - b_T(s_i(1 - d, \vec{r}))}{b_T(s_0(\vec{r}))}$$

for $d \in [0, D]$. This shows that $q \cdot F_p(T, \delta) \geq 1$ is a necessary condition for the seller to be always truthful; in particular, the condition is necessary and sufficient to ensure truthfulness when the current score of the seller is $s_0(\vec{r}) = 1$. To show that it is also sufficient we will show that

$$\frac{\sum_{i=0}^{T-1} \delta^i (b_T(1) - b_T(1 - d/T))}{b_T(1)} \leq \frac{\sum_{i=0}^{T-1} \delta^i (b_T(s_i(1, \vec{r})) - b_T(s_i(1, \vec{r}) - d/T))}{b_T(s_0(\vec{r}))}$$

for all \vec{r} and $d > 0$. Since $s_i(1, \vec{r}) \geq s_0(\vec{r})$, it suffices that

$$(b_T(1) - b_T(1-t))b_T(k) \leq b_T(1)(b_T(k) - b_T(k-t))$$

for all $k < 1$ and $t \in (0, k]$, which holds (by Lemma 2).

Finally, we observe that the function F_p is increasing in δ . ■

Proof of Proposition 2: We will show that $\log(F_p(T, \delta))$ satisfies increasing differences. Let $T' \geq T$.

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

where $t(T, T')$ does not depend on δ . Thus, to show that $\log(F_p(T, \delta))$ 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 $\log(F_p(T, \delta))$ has increasing differences in (T, δ) ; the result follows by applying Topkis' Theorem (Topkis 1998). ■

Proof of Proposition 3: $R(q, T, \delta)$ is the minimum ratio (over all rating histories and all possible deviations) of future decrease in payoff to current gain in payoff, if the seller deviates from being

truthful in the current period. In other words, it is optimal for the seller to be always truthful if and only if $R(q, T, \delta) \geq 1$ (by the definition of $R(q, T, \delta)$). Thus, to show that $R(q, T, \delta) = q \cdot F_m(T, \delta)$ it suffices to show that it is optimal for the seller to be always truthful if and only if $q \cdot F_m(T, \delta) \geq 1$. We proceed by showing the latter in this proof.

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. The seller may consider to exaggerate an item in his description if in that period he has a low value item for sale. Let $\hat{V}(\vec{r})$ be the expected infinite horizon discounted payoff to the seller if he is always truthful and his current vector of ratings is \vec{r} . The seller does not deviate from the truthful policy at \vec{r} if and only if

$$(v_H - v_L) \cdot b_T \left(\sum_{i=1}^T r_i / T \right) \leq \delta [(1 - p_{tL}) \hat{V}(1, \vec{r}) - (1 - p_{uL}) \hat{V}(0, \vec{r}) + (p_{tL} - p_{uL}) \hat{V}(\vec{r})]. \quad (5)$$

With probability $a(i, j) \equiv \binom{i}{j} (1-p)^j p^{i-j}$ exactly j new ratings arrive in i periods. Let $s_j(\vec{r})$ be the score of the seller after j good rating arrive with initial rating vector \vec{r} . Then $s_0(\vec{r}) = \sum_{i=1}^T r_i / T$, and

$$\begin{aligned} & (1 - p_{tL}) \hat{V}(1, \vec{r}) - (1 - p_{uL}) \hat{V}(0, \vec{r}) + (p_{tL} - p_{uL}) \hat{V}(\vec{r}) = \\ & q \cdot \sum_{i=0}^{\infty} \delta^i \sum_{j=0}^{\min(T-1, i)} a(i, j) [(1 - p_{tL}) b_T(s_j(1, \vec{r})) - (1 - p_{uL}) b_T(s_j(0, \vec{r})) + (p_{tL} - p_{uL}) b_T(s_j(\vec{r}))] = \\ & q \cdot \sum_{i=0}^{\infty} \delta^i \sum_{j=0}^{\min(T-1, i)} a(i, j) [(1 - p_{tL}) (b_T(s_j(1, \vec{r})) - b_T(s_j(\vec{r}))) + (1 - p_{uL}) (b_T(s_j(\vec{r})) - b_T(s_j(0, \vec{r})))] \end{aligned}$$

This shows that $q \cdot F_m(T, \delta) \geq 1$ is equivalent to (5) for $\vec{r} = \vec{1}$, and thus it is a necessary condition for the seller to be always truthful. To show sufficiency, it suffices to show that for $j < T$,

$$\begin{aligned} & \frac{(1 - p_{tL}) (b_T(s_j(1, \vec{r})) - b_T(s_j(\vec{r}))) + (1 - p_{uL}) (b_T(s_j(\vec{r})) - b_T(s_j(0, \vec{r})))}{b_T(s_0(\vec{r}))} \geq \\ & \frac{(1 - p_{uL}) (b_T(1) - b_T(1 - 1/T))}{b_T(1)} \end{aligned}$$

Note that $s_j(1, \vec{r}) \geq s_j(\vec{r}) \geq s_j(0, \vec{r})$ and $s_j(1, \vec{r}) = s_j(0, \vec{r}) + 1/T$. Therefore, either $s_j(1, \vec{r}) = s_j(\vec{r}) + 1/T$ or $s_j(\vec{r}) = s_j(0, \vec{r}) + 1/T$.

We first assume that $s_j(1, \vec{r}) = s_j(\vec{r}) + 1/T$ and $s_j(\vec{r}) = s_j(0, \vec{r})$. Then

$$\begin{aligned} & \frac{(1 - p_{tL})(b_T(s_j(1, \vec{r})) - b_T(s_j(\vec{r}))) + (1 - p_{uL})(b_T(s_j(\vec{r})) - b_T(s_j(0, \vec{r})))}{b_T(s_0(\vec{r}))} = \\ & (1 - p_{tL}) \frac{b_T(s_j(1, \vec{r})) - b_T(s_j(\vec{r}))}{b_T(s_0(\vec{r}))} \geq \\ & (1 - p_{tL}) \frac{b_T(1) - b_T(1 - 1/T)}{b_T(1)} \geq \\ & (1 - p_{uL}) \frac{b_T(1) - b_T(1 - 1/T)}{b_T(1)} \end{aligned}$$

This first inequality is a consequence of the fact that b_T is logarithmically concave, together with Lemma 2; and the second inequality holds because $p_{uL} \geq p_{tL}$.

We next assume that $s_j(1, \vec{r}) = s_j(\vec{r})$ and $s_j(\vec{r}) = s_j(0, \vec{r}) + 1/T$. Then

$$\begin{aligned} & \frac{(1 - p_{tL})(b_T(s_j(1, \vec{r})) - b_T(s_j(\vec{r}))) + (1 - p_{uL})(b_T(s_j(\vec{r})) - b_T(s_j(0, \vec{r})))}{b_T(s_0(\vec{r}))} = \\ & (1 - p_{uL}) \frac{b_T(s_j(\vec{r})) - b_T(s_j(0, \vec{r}))}{b_T(s_0(\vec{r}))} \geq \\ & (1 - p_{uL}) \frac{b_T(1) - b_T(1 - 1/T)}{b_T(1)} \end{aligned}$$

This again holds by applying (4) of Lemma 2.

We conclude that if $p_{uL} \geq p_{tL}$, then the seller is always truthful if and only if $q \cdot F_m(T, \delta) \geq 1$.

Finally, we observe that the function F_m is increasing in δ . ■

Proof of Proposition 4: Let

$$\begin{aligned} g(T, \delta) & \equiv \frac{b_T(1) - b_T(1 - 1/T)}{b_T(1)} \sum_{i=0}^{\infty} f(T, i) \delta^i, \\ f(T, i) & = \sum_{j=0}^{\min(T-1, i)} \binom{i}{j} (1-p)^j p^{i-j}. \end{aligned}$$

Also let $a(i, j) \equiv \binom{i}{j} (1-p)^j p^{i-j}$. Clearly, $T_m^*(\delta) = \arg \max_T \{g(T, \delta)\}$.

We will show that g satisfies the single crossing property in (T, δ) , i.e., that $g(T', \delta) > g(T, \delta)$ implies $g(T', \delta') > g(T, \delta')$ for $\delta' > \delta$ and $T' > T$. This will imply that $T^*(\delta)$ is increasing in δ (Milgrom and Shannon 1994). Let

$$c \equiv \frac{(b_{T'}(1) - b_{T'}(1 - 1/T'))/b_{T'}(1)}{(b_T(1) - b_T(1 - 1/T))/b_T(1)}.$$

Equivalently we will show that if

$$\sum_{i=0}^{\infty} (c \cdot f(T', i)) \delta^i > \sum_{i=0}^{\infty} f(T, i) \delta^i,$$

then the inequality also holds for $\delta' > \delta$.

We first show that if $c \cdot f(T', i) > f(T, i)$, then $c \cdot f(T', i+1) > f(T, i+1)$. We consider the cases $i < T$ and $i \geq T$ separately.

Suppose $i < T$. Then $c f(T', i) > f(T, i)$ implies that

$$c \sum_{j=0}^i a(i, j) > \sum_{j=0}^i a(i, j),$$

which can only happen if $c > 1$. Also $\min(i+1, T' - 1) = i+1$, while $\min(i+1, T - 1)$ is $i+1$ if $i < T - 1$; and i if $i = T - 1$. In either case, $c \cdot f(T', i+1) > f(T, i+1)$.

Now suppose that $i \geq T$, and let $k \equiv \min(i, T' - 1)$. Assume $c \cdot f(T', i) - f(T, i) > 0$. Then

$$\sum_{j=T}^k a(i, j) > \frac{1-c}{c} \sum_{j=0}^{T-1} a(i, j).$$

We observe that

$$a(i+1, j) = (1-p) \frac{i+1}{i+1-j} a(i, j) = (1-p) \left(1 + \frac{j}{i+1-j} \right) a(i, j).$$

Then,

$$\begin{aligned} & \sum_{j=T}^k a(i+1, j) - \frac{1-c}{c} \sum_{j=0}^{T-1} a(i+1, j) \\ &= (1-p) \sum_{j=T}^k \left(1 + \frac{j}{i+1-j} \right) a(i, j) - (1-p) \frac{1-c}{c} \sum_{j=0}^{T-1} \left(1 + \frac{j}{i+1-j} \right) a(i, j) \end{aligned}$$

Moreover,

$$\sum_{j=T}^k \frac{i+1}{i+1-j} a(i, j) \geq \frac{i+1}{i+1-T} \sum_{j=T}^k a(i, j) > \frac{1-c}{c} \frac{i+1}{i+1-T} \sum_{j=0}^{T-1} a(i, j) \geq \frac{1-c}{c} \sum_{j=0}^{T-1} \frac{i+1}{i+1-j} a(i, j).$$

Since

$$\sum_{j=T}^k a(i, j) - \frac{1-c}{c} \sum_{j=0}^{T-1} a(i, j) > 0,$$

we have that

$$\sum_{j=T}^k \frac{i+1}{i+1-j} a(i, j) \geq \frac{i+1}{i+1-T} \sum_{j=T}^k a(i, j) > \frac{1-c}{c} \frac{i+1}{i+1-T} \sum_{j=0}^{T-1} a(i, j) \geq \frac{1-c}{c} \sum_{j=0}^{T-1} \frac{i+1}{i+1-j} a(i, j)$$

We conclude that if $c \cdot f(T', i) - f(T, i) > 0$, then

$$\sum_{j=T}^k a(i+1, j) > \frac{1-c}{c} \sum_{j=0}^{T-1} a(i+1, j).$$

Since $a(i+1, j) \geq 0$, $\min(i+1, T'-1) \geq k$ and $\min(i+1, T-1) = T-1$, this implies that

$$c \cdot f(T', i+1) - f(T, i+1) > 0.$$

The final step of the proof is to show that if $\sum_{i=0}^{\infty} (c \cdot f(T', i)) \delta^i > \sum_{i=0}^{\infty} f(T, i) \delta^i$, then the inequality also holds for $\delta' > \delta$. Let $T' > T$ and $e_i = c \cdot f(T', i) - f(T, i)$. We have shown that if $e_i > 0$ then $e_{i+1} > 0$. If $e_i > 0$ for all i , then trivially $\sum_{i=0}^{\infty} e_i \delta^i > 0$ for all δ .

Now suppose $e_0 < 0$ and let $k = \max\{i : e_i < 0\}$. If $\sum_{i=0}^{\infty} e_i \delta^i > 0$, then

$$\sum_{i=k+1}^{\infty} |e_i| \delta^i > \sum_{i=1}^k |e_i| \delta^i.$$

Moreover,

$$\sum_{i=k+1}^{\infty} i |e_i| \delta^{i-1} > \sum_{i=1}^k i |e_i| \delta^{i-1}.$$

The last inequality is equivalent to

$$\frac{\partial}{\partial \delta} \left(\sum_{i=0}^{\infty} e_i \delta^i \right) = \sum_{i=1}^{\infty} i e_i \delta^{i-1} > 0,$$

which concludes the proof. ■

Proof of Proposition 5: We first observe that the function F_w is increasing in δ .

$R(q, T, \delta)$ is the minimum ratio (over all rating histories and all possible deviations) of future decrease in payoff to current gain in payoff, if the seller deviates from being truthful in the current period. In other words, it is optimal for the seller to be always truthful if and only if $R(q, T, \delta) \geq 1$ (by the definition of $R(q, T, \delta)$). Thus, to show that $R(q, T, \delta) = q \cdot F_w(T, \delta)$ it suffices to show that it is optimal for the seller to be always truthful if and only if $q \cdot F_w(T, \delta) \geq 1$. We proceed by showing the latter in this proof.

It is optimal for the seller to describe a low value item truthfully at \vec{r} if

$$(v_H - v_L) \cdot b_T \left(\sum_{i=1}^T r_i / T \right) \leq \delta (1 - p_{uL} - p_{tL}) (V(1, \vec{r}) - V(0, \vec{r})).$$

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. The seller may consider to exaggerate an item in his description if in that period he has a low value item for sale. Let $\hat{V}(\vec{r})$ be the infinite horizon discounted payoff to the seller if he is always truthful. The seller will not deviate from being truthful when his ratings are \vec{r} if

$$(v_H - v_L) \cdot b_T \left(\sum_{j=1}^T r_j / T \right) \leq \delta(1 - p_{uL} - p_{tL})(\hat{V}(1, \vec{r}) - \hat{V}(0, \vec{r})).$$

We observe that the payments to the seller from $\hat{V}(1, \vec{r})$ and $\hat{V}(0, \vec{r})$ may differ in the next T periods, but not after that. Let $s_i(\vec{r})$ be the seller's score in i periods if his current rating vector is \vec{r} and he gets good ratings in all future periods. Let $a(i, k) \equiv \binom{i}{k} p^k (1-p)^{i-k}$; this is the probability that the seller gets k bad ratings (which are inaccurate) after i periods. Then,

$$\hat{V}(1, \vec{r}) - \hat{V}(0, \vec{r}) = q \sum_{i=0}^{T-1} \sum_{k=0}^i a(i, k) (b_T(s_i(1, \vec{r}) - k/T) - b_T(s_i(1, \vec{r}) - (k+1)/T))$$

We conclude that the seller will not deviate from being truthful when his ratings are \vec{r} if

$$(v_H - v_L) \cdot b_T(s_0(\vec{r})) \leq \delta(1 - p_{uL} - p_{tL}) q \sum_{i=0}^{T-1} \sum_{k=0}^i a(i, k) (b_T(s_i(1, \vec{r}) - k/T) - b_T(s_i(1, \vec{r}) - (k+1)/T)).$$

Let

$$f(i, k) \equiv \frac{b_T(s_i(1, \vec{r}) - k/T) - b_T(s_i(1, \vec{r}) - (k+1)/T)}{b_T(s_0(\vec{r}))} - \frac{b_T(1 - k/T) - b_T(1 - (k+1)/T)}{b_T(1)}$$

and

$$g(i, k) \equiv \frac{b_T(1 - k/T)}{b_T(1)} - \frac{b_T(s_i(1, \vec{r}) - k/T)}{b_T(s_i(1, \vec{r}))}.$$

To prove the proposition it suffices to show that

$$\sum_{k=0}^i a(i, k) f(i, k) \geq 0. \quad (6)$$

We have that

$$\begin{aligned} f(i, k) &= \frac{b_T(s_i(1, \vec{r}) - k/T) - b_T(s_i(1, \vec{r}) - (k+1)/T)}{b_T(s_0(\vec{r}))} - \frac{b_T(1 - k/T) - b_T(1 - (k+1)/T)}{b_T(1)} \\ &\geq \frac{b_T(s_i(1, \vec{r}) - k/T) - b_T(s_i(1, \vec{r}) - (k+1)/T)}{b_T(s_i(1, \vec{r}))} - \frac{b_T(1 - k/T) - b_T(1 - (k+1)/T)}{b_T(1)} \end{aligned}$$

$$\begin{aligned}
&= \frac{b_T(1 - (k+1)/T)}{b_T(1)} - \frac{b_T(s_i(1, \vec{r}) - (k+1)/T)}{b_T(s_i(1, \vec{r}))} + \frac{b_T(s_i(1, \vec{r}) - k/T)}{b_T(s_i(1, \vec{r}))} - \frac{b_T(1 - k/T)}{b_T(1)} \\
&= g(i, k+1) - g(i, k)
\end{aligned}$$

Moreover, there exists λ such that

$$\frac{b_T(x - k/T) - b_T(x - (k+1)/T)}{b_T(x)} - \frac{b_T(1 - k/T) - b_T(1 - (k+1)/T)}{b_T(1)} \geq -\lambda$$

for $x \in \{1/T, 2/T, \dots, (T-1)/T\}$ and $k \in \{1, 2, \dots, Tx-1\}$. For instance, the aforementioned inequality holds for any premium function b_T if $\lambda = 2$, since

$$\frac{b_T(x - (k+1)/T)}{b_T(x)} + \frac{b_T(1 - k/T)}{b_T(1)} \leq \frac{b_T(x)}{b_T(x)} + \frac{b_T(1)}{b_T(1)} = 2.$$

Since b_T is logarithmically concave, the function $\log(b_T(x)) - \log(b_T(x-w))$ is decreasing in x .

Thus $b_T(x-w)/b_T(x)$ is increasing in s and $g(i, k) \geq 0$. We consider three cases:

- Case 1: $g(i, 1) > 0$

Then $f(i, 0) > 0$ and condition 6 holds for

$$p \leq 1 - \left(\frac{\lambda}{f(i, 0) + \lambda} \right)^{1/i}.$$

- Case 2: $g(i, k) = 0$ for $k \in \{0, \dots, i+1\}$

Then condition 6 trivially holds.

- Case 3: There exists $k^* \in \{1, \dots, i\}$ such that $g(i, k) = 0$ for $k \leq k^*$ and $g(i, k^* + 1) > 0$

Then $f(i, k^*) > 0$ and $f(i, j) \geq -$ for $k < k^*$. We conclude that condition (6) holds for

$$p \leq 1 - \left(\frac{\lambda}{f(i, k^*) + \lambda} \right)^{1/i}.$$

We next show a more general result than the statement of the proposition: Property 1 holds with inaccurate ratings for arbitrary values of p if

- (i) b_T is concave; or
- (ii) b_T is an exponential function, i.e., $b_T(s) = e^{\alpha \cdot s + \beta}$ with $\alpha > 0$.

The remainder of the proof shows that if (i) or (ii) is satisfied, then (6) holds.

We first consider condition (i) and assume that b_T is concave. Then,

$$b_T(s_i(1, \vec{r}) - k/T) - b_T(s_i(1, \vec{r}) - (k+1)/T) \geq b_T(1 - k/T) - b_T(1 - (k+1)/T)$$

by the concavity of b_T , and

$$b_T(1) \geq b_T\left(\sum_{j=1}^T r_j/T\right)$$

since b_T is increasing. We conclude that in this case

$$\left(\frac{b_T(s_i(1, \vec{r}) - k/T) - b_T(s_i(1, \vec{r}) - (k+1)/T)}{b_T\left(\sum_{i=1}^T r_i/T\right)} - \frac{b_T(1 - k/T) - b_T(1 - (k+1)/T)}{b_T(1)}\right) \geq 0,$$

and thus (6) holds.

Now assume that $b_T(s) = e^{\alpha \cdot s + \beta}$ and $\alpha > 0$ (condition (ii)). Then

$$\begin{aligned} & \frac{b_T(s_i(1, \vec{r}) - k/T) - b_T(s_i(1, \vec{r}) - (k+1)/T)}{b_T\left(\sum_{i=1}^T r_i/T\right)} - \frac{b_T(1 - k/T) - b_T(1 - (k+1)/T)}{b_T(1)} = \\ & \frac{e^{\alpha(s_i(1, \vec{r}) - k/T) + \beta} - e^{\alpha(s_i(1, \vec{r}) - (k+1)/T) + \beta}}{e^{\alpha\left(\sum_{j=1}^T r_j/T\right) + \beta}} - \frac{e^{\alpha(1 - k/T) + \beta} - e^{\alpha(1 - (k+1)/T) + \beta}}{e^{\alpha + \beta}} = \\ & \frac{e^{\alpha(s_i(1, \vec{r}) - k/T)} - e^{\alpha(s_i(1, \vec{r}) - (k+1)/T)}}{e^{\alpha\left(\sum_{j=1}^T r_j/T\right)}} - \frac{e^{\alpha(1 - k/T)} - e^{\alpha(1 - (k+1)/T)}}{e^{\alpha}} = \\ & \left(\frac{e^{\alpha(s_i(1, \vec{r}))}}{e^{\alpha\left(\sum_{j=1}^T r_j/T\right)}} - 1\right) (e^{-\alpha k/T} - e^{-\alpha(k+1)/T}) \geq 0 \end{aligned}$$

because $s_i(1, \vec{r}) \geq \sum_{j=1}^T r_j/T$ and $\alpha > 0$. This implies that (6) holds. ■

Proof of Proposition 6: Let

$$\begin{aligned} c(i) &= \sum_{k=0}^i \binom{i}{k} p^k (1-p)^{i-k} f(k), \\ f(k) &= \frac{b_T(1 - k/T) - b_T(1 - 1/T - k/T)}{b_T(1)} - \frac{b_{T'}(1 - k/T') - b_{T'}(1 - 1/T' - k/T')}{b_{T'}(1)}, \\ g(T, \delta) &= \sum_{i=0}^{T-1} \delta^i \sum_{k=0}^i \binom{i}{k} p^k (1-p)^{i-k} \frac{b_T(1 - k/T) - b_T(1 - (k+1)/T)}{b_T(1)}. \end{aligned}$$

Clearly, $T_w^*(\delta) = \arg \max_T \{g(T, \delta)\}$.

The proof consists of three steps. First, we show that if $h_{T, T'}$ satisfies the assumption of the proposition, then $f(k) < 0$ implies $f(k+1) < 0$. The second step is to show that $c(i) < 0$ implies that $c(i+1) < 0$. Then we show that g satisfies the single crossing property in (T, δ) and conclude that $T_w^*(\delta)$ is increasing.

Step 1: Let

$$h(k) = \frac{b_{T'}(1 - k/T')}{b_{T'}(1)} - \frac{b_T(1 - k/T)}{b_T(1)}$$

We observe that

$$f(k) = \frac{b_{T'}(1 - (k+1)/T')}{b_{T'}(1)} - \frac{b_T(1 - (k+1)/T)}{b_T(1)} - \frac{b_{T'}(1 - k/T')}{b_{T'}(1)} + \frac{b_T(1 - k/T)}{b_T(1)} = h(k+1) - h(k).$$

Since there exists a $k_0 \in \{0, \dots, T\}$ such that h_k is increasing in k for $k < k_0$ and decreasing in k for $k \geq k_0$, we conclude that $f_k < 0$ if and only if $k \geq k_0$. Thus, if $f(k) < 0$ then $f(k+1) < 0$.

Step 2: Let

$$a(i, k) = \binom{i}{k} p^k (1-p)^{i-k}.$$

The key property we exploit is that:

$$\frac{a(i+1, k)}{a(i, k)} = (1-p) \frac{i+1}{i+1-k}$$

is strictly increasing in k . We have shown that $f(k) \geq 0$ for $k < k_0$ and $f(k) < 0$ for $k \geq k_0$. Suppose $c(i) < 0$. Of course, in this case we must have $i \geq k_0$. Then:

$$\sum_{k=0}^{k_0-1} a(i, k) f(k) < - \sum_{k=k_0}^i a(i, k) f(k).$$

But now note that for all $k < k_0$, $i+1-k > i+1-k_0$; and for all k such that $k_0 \leq k < i$, $i+1-k \leq i+1-k_0$. So we get:

$$\begin{aligned} \sum_{k=0}^{k_0-1} a(i+1, k) f(k) &= (i+1)(1-p) \sum_{k=0}^{k_0-1} a(i, k) f(k) / (i+1-k) \\ &< (i+1)(1-p) \sum_{k=0}^{k_0-1} a(i, k) f(k) / (i+1-k_0) \\ &< -(i+1)(1-p) \sum_{k=k_0}^i a(i, k) f(k) / (i+1-k_0) \\ &< -(i+1)(1-p) \sum_{k=k_0}^i a(i, k) f(k) / (i+1-k) \\ &< - \sum_{k=k_0}^{i+1} a(i+1, k) f(k) \end{aligned}$$

where the last inequality follows since $f(i+1) < 0$. We conclude that $c(i+1) < 0$, as required.

Step 3: Let $T' > T$ and $\delta' > \delta$. The function g satisfies the single crossing property in (T, δ) if $g(T', \delta) > g(T, \delta)$ implies that $g(T', \delta') > g(T, \delta')$. We observe that

$$g(T', x) - g(T, x) =$$

$$\begin{aligned} & \sum_{i=T}^{T'-1} \delta^i \sum_{k=0}^i \binom{i}{k} p^k (1-p)^{i-k} \frac{b_T(1-k/T') - b_T(1-(k+1)/T')}{b_{T'}(1)} - \sum_{i=0}^{T-1} \delta^i \sum_{k=0}^i \binom{i}{k} p^k (1-p)^{i-k} f(k) = \\ & \sum_{i=T}^{T'-1} \delta^i \sum_{k=0}^i \binom{i}{k} p^k (1-p)^{i-k} \frac{b_T(1-k/T) - b_T(1-(k+1)/T)}{b_T(1)} - \sum_{i=0}^{T-1} \delta^i c(i) \end{aligned}$$

According to Step 2, there exists some i_0 such that we can rewrite the previous difference as

$$g(T', x) - g(T, x) = - \sum_{i=0}^{i_0-1} x^i d_i + \sum_{i=i_0}^{T'-1} x^i d_i,$$

where $d_i \geq 0$ for all i . Assume that $g(T', \delta) - g(T, \delta) > 0$. Then,

$$\sum_{i=0}^{i_0-1} i \delta^{i-1} d_i = \sum_{i=0}^{i_0-1} \frac{i}{\delta} \delta^i d_i \leq \frac{i_0-1}{\delta} \sum_{i=0}^{i_0-1} \delta^i d_i \leq \frac{i_0-1}{\delta} \sum_{i=i_0}^{T'-1} \delta^i d_i = \sum_{i=T}^{T'-1} (i_0-1) \delta^{i-1} d_i \leq \sum_{i=i_0}^{T'-1} i \delta^{i-1} d_i.$$

This implies that if $g(T', \delta) - g(T, \delta) > 0$, then $g'(T', \delta) - g'(T, \delta) \geq 0$. We conclude that if $g(T', \delta) - g(T, \delta) > 0$, then $g(T', \delta') - g(T, \delta') > 0$ for $\delta' > \delta$. This shows that the objective satisfies the single crossing property. Thus, we apply Theorem 4 from Milgrom and Shannon (1994) to conclude the proof. \blacksquare

Proof of Corollary 1: We will first show a stronger result: if $b'(1-y) - yb''(1-y) < 0$ for some $y \in [0, 1]$ implies that $b'(1-z) - zb''(1-z) < 0$ for $z > y$, then $T_w^*(\delta)$ is increasing in δ . Then, we will show that this is the case if $\log(b'(s))$ is concave.

Let

$$g(x) \equiv b(1) \cdot h'_{T',T}(x) = \frac{1}{T} b'(1-x/T) - \frac{1}{T'} b'(1-x/T').$$

Clearly, $g(0) > 0$. It suffices to show that if $g(x) < 0$ then $g(x') < 0$ for $x' > x$ in order to apply Proposition 6.

$$\begin{aligned} g(x) &= \frac{1}{T} b'(1-x/T) - \frac{1}{T'} b'(1-x/T') \\ &= \int_{1/T'}^{1/T} \frac{\partial}{\partial y} [y b'(1-yx)] dy \\ &= \int_{1/T'}^{1/T} [b'(1-yx) - yx b''(1-yx)] dy \\ &= \frac{1}{x} \int_{x/T'}^{x/T} [b'(1-z) - z b''(1-z)] dz \end{aligned}$$

If $b'(1-y) - yb''(1-y) > 0$ for all $y \in [0, 1]$ then $g(x) > 0$ for all $x \in [0, T]$. Otherwise there exists $z_0 \in (0, 1]$ such that $b'(1-z_0) - z_0 b''(1-z_0) = 0$ and $b'(1-z) - z b''(1-z) > 0$ for $z < z_0$; $b'(1-z) - z b''(1-z) < 0$ for $z > z_0$.

If $g(x) < 0$, then

$$\int_{x/T'}^{z_0} [b'(1-z) - zb''(1-z)]dz < \int_{z_0}^{x/T} |b'(1-z) - zb''(1-z)|dz.$$

Let $x' > x$. Then $x'/T' > x/T'$ and $x'/T > x/T$. If $x'/T' > z_0$, then $g(x') < 0$. If $x'/T' < z_0$, then

$$\int_{x'/T'}^{z_0} [b'(1-z) - zb''(1-z)]dz > \int_{x'/T'}^{z_0} [b'(1-z) - zb''(1-z)]dz$$

and

$$\int_{z_0}^{x/T} |b'(1-z) - zb''(1-z)|dz < \int_{z_0}^{x'/T} |b'(1-z) - zb''(1-z)|dz.$$

We conclude that

$$\int_{x'/T'}^{z_0} [b'(1-z) - zb''(1-z)]dz < \int_{z_0}^{x'/T} |b'(1-z) - zb''(1-z)|dz$$

which implies that $g(x') < 0$.

We have shown the result for the case that $b'(1-y) - yb''(1-y) < 0$ implies that $b'(1-z) - zb''(1-z) < 0$ for $z > y$. A sufficient condition for this is that $(1-x)b''(x)/b'(x)$ is decreasing. The latter holds if b' is logarithmically concave. ■