Welfare-Improving Cascades and the Effect of Noisy Reviews

Nick Arnosti* and Daniel Russo**

Stanford University

Abstract. We study a setting in which firms produce items whose quality is ex-ante unobservable, but learned by customers over time. Firms take customer learning into account when making production decisions. We focus on the effect that the review process has on product quality. Specifically, we compare equilibrium quality levels in the setting described above to the quality that would be produced if customers could observe item quality directly. We find that in many cases, customers are better off when relying on reviews, i.e. better off in the world where they have less information. The idea behind our result is that the risk of losing future profits due to bad initial reviews may drive firms to produce an exceptional product. This intuitive insight contrasts sharply with much of the previous academic literature on the subject.

1 Introduction

It is often impossible to directly determine the quality of an item before buying it. When this is the case, potential customers try to learn from the experiences of others. Increasingly, they rely on online reviews to help with their decisions. These reviews can significantly influence a firm's profitability: a Harvard Business School study [11] recently concluded that each additional star on Yelp generates (on average) a 5-9% increase in revenue for small businesses. Not only does firm success depend on reviews, but evidence is mounting that business owners are *aware* of this, and take it into account when making decisions. Some businesses respond directly to customers who leave negative reviews, and try to make amends. Others change their business practices: one Chicago bookstore "totally revamped our customer service approach" due to reviews left on Yelp [17].

This paper studies a setting in which customers learn from reviews left by others, and firms make choices with this fact in mind. There are many interesting questions that one could ask about such a market. We focus on the quality of items produced. Intuition suggests that reviews are an imperfect substitute for directly observing item quality. Formalizing this idea, the literature on this topic almost universally reaches the conclusion that product quality is highest when

* Supported in part by DARPA under the GRAPHS program.

** Supported by a Burt and Deedee McMurty Stanford Graduate Fellowship.

Y. Chen and N. Immorlica (Eds.): WINE 2013, LNCS 8289, pp. 405–420, 2013.

customers can observe it. We illustrate that when item quality is endogenously chosen by firms, and firms take customer learning into account, the *opposite* might be true.

In our model, firms begin by making an irreversible decision about the quality of items that they will produce. Customers cannot observe this decision, but whenever a customer patronizes the firm they leave behind a review. Each review provides a noisy signal of the firms' item quality, which future customers use to draw inferences. Even high quality firms are at risk of losing business due to a bad review. To guard against this possibility, firms may produce a product that is better than customers would demand if they could directly observe its quality.

We clarify this intuition through two stylized models. In each model, firms are distinguished by their cost structures: some firms can provide high quality items more cheaply than others. We first consider a one-shot model, in which customer decisions are based on a single review. We then extend this to an infinite horizon model in which firms are visited by a sequence of customers. Customers only purchase from firms with a sufficiently favorable review history, and the only way that a firm can signal its quality is through customer reviews. This can lead to a cascade-like phenomenon: if early customers leave bad reviews, later ones may choose not to purchase. When this occurs, customers stop learning about the firm, and so the firm may go out of business even if they are producing high quality goods. The threat of this cascade gives firms an incentive to set higher initial quality levels.

The remainder of the paper is organized as follows. Section 2 places our work in the context of related academic literature. We introduce and discuss our general framework in Section 3. From there, we describe and analyze a single-period model in Section 4 before considering the infinite-horizon case in Section 5. We close with a summary of our results and a discussion of possibilities for future work.

2 Literature Review

The topic of signaling and reputation in markets has a rich history in economics. In a groundbreaking paper, Spence [16] introduces a simple model of job market signaling. In his model, intrinsically high quality workers put in effort in school in order to signal their ability to potential employers and, by doing so, differentiate themselves from low quality workers for whom sending positive signals is more costly. The firms in our model could be naturally re-interpreted as workers attempting to maintain a good reputation, and the customers as potential employers deciding whether to hire the worker. Seen in this light, our work differs from Spence's in at least two ways. The first is that we model the observation of individual performance as a noisy process. The second is that the effort employees exert in order to signal their quality directly benefits their employers.

In our model a potentially long-lived firm interacts with a series of short-lived customers. Thus, though we speak of a firm's "reputation," our setting differs substantially from the large literature that studies the role of reputation in long run strategic relationships (see for example Rubinstein [13], Fudenberg and Maskin [7], Fudenberg et al. [8]).

One distinctive component of our model is that firms receive a review only if a customer purchases their product. Because of this, a firm could be unable to repair its reputation after receiving several negative reviews, since it must attract customers in order to signal its quality. In this way, our work relates to the study of social herding and information cascades, as presented in the seminal work of Banerjee [2] and Bikhchandani et al. [4]. The main idea in these papers is that when customers have imperfect information and make decisions sequentially, it may be rational for them to ignore their private information and instead mimic the actions of those who went before them. This can induce a "cascade" in which each later customer takes the same action, even if it is in fact a poor one. The rather pessimistic message from these papers is that the outcome that results from a sequence of individually rational decisions may be arbitrarily worse than what results from a socially optimal decision rule.

There are effectively two types of information cascade: those in which customers repeatedly patronize a firm that is producing items of disappointingly low quality, and those in which customers abandon a firm producing high quality items. In our model, the former cannot happen, but the latter is a possibility. Rather than lament this inefficiency, we take a much cheerier perspective. In particular, our results indicate that the threat of such a cascade may cause firms to set higher quality levels than they otherwise would. Viewed in this light, when the underlying state of the world (i.e. product quality in our model) is not exogenously determined but rather strategically selected, the possibility of herd behavior may actually *enhance* consumer welfare.

Even closer to our work is a set of papers that deal specifically with models in which customers must decide whether to purchase items whose quality they cannot observe. Examples include Smallwood and Conlisk [15], Shapiro [14], Rogerson [12], Allen [1], Wolinsky [18], Cooper and Ross [6], Hörner [9], Bar-Isaac [3], Bose et al. [5], Ifrach et al. [10]. These models differ on a number of dimensions, such as the information possessed by firms and customers, whether customers behave strategically, and the role of item price.

Both Wolinsky [18] and Allen [1] construct equilibria where price signals quality perfectly to consumers. While these may help us understand some markets, they preclude the study of learning effects. At the other end of the spectrum, Cooper and Ross [6], Hörner [9] and Bar-Isaac [3] construct equilibria where posted prices depend only on information already available to consumers, i.e. the history of product reviews. Although prices affect firm incentives in these models, customers do not use them for inference.

These examples highlight the difficulty of building models where price plays an interesting role in the customer learning process. If firms are informed, consumers are rational, and firms can set prices, equilibrium prices are typically either fully revealing or do not reveal any new information to customers. To see why this is true, consider pure strategy equilibria when there are only two firm types. If prices set by the two types of firms depend on information that is not available to customers, then so long as the prices are not equal, customers can use them to infer firm type. Because we wish to focus on *learning* in reputation games, we choose not to incorporate price directly into our model. In this sense, our work applies to markets in which prices do not notably differentiate products and customers must look elsewhere in order to infer product quality.

Bar-Isaac [3] and Bose et al. [5] avoid the problem of firm prices revealing quality by studying models in which firms have no more information than consumers do. Alternatively, Smallwood and Conlisk [15] and Ifrach et al. [10] do not model customers as strategic. Instead, they exogenously specify consumer behavior, and then study the market trajectory that results from various firm choices. Ifrach et al. [10] argue that assuming rationality "introduces a formidable analytical and computational onus on each agent that may be hard to justify as a model of actual choice behavior." While this is no doubt true in many cases, a strength of this paper is that customer best response dynamics in our model turn out to be quite simple to describe.

Another critical feature of our model is that, like Rogerson [12] and Shapiro [14], we study a scenario where firms choose product quality at the beginning of the game, and it remains fixed throughout. This model is appropriate for industries in which product quality is derived from high-cost training or long-term investment in capital. Other industries would be more appropriately modeled by allowing firms to choose their quality level dynamically. There are a number of interesting questions in this case, which we discuss briefly towards the end of the paper.

Although the papers discussed above differ in many respects, they agree on at least one point, expressed succinctly by Smallwood and Conlisk [15]: "consumers pay considerably for being ill-informed." Wolinsky [18] finds that as the signal to customers becomes more informative, equilibrium prices drop. In Rogerson [12], a more informative signal results in more high quality goods in the market. Allen [1] finds that making quality unobservable may not change equilibrium behavior, but if it does, it either results in a lower quantity being offered at a higher price than before, or collapses the market entirely. Cooper and Ross [6] observe that when some customers cannot determine product quality, the price offered is the same as under full information, but the average quality in the market is lower.

Shapiro [14] reports that any "self-fulfilling quality level" must lie below the quality of goods produced in the complete information case, and that as information about the firm's quality spreads more rapidly (i.e. approaching full information), the self-fulfilling quality level rises. Though there are many differences between Shapiro's model and the one we consider here, the most important is that in his model, reputation evolves deterministically, gradually shifting from the customer's initial expectations to the true underlying quality. Stochasticity is the key to our result that firms may produce higher quality goods when quality cannot be directly monitored: even firms producing items that exceed customer expectations have reason to fear a bad review.

3 Model Introduction

3.1 Game Rules

We consider the following general framework. Each firm has a type T, which can be either high, H, or low, L. Customers cannot observe a firm's type, but it is common knowledge that each firm is high-type with probability $\alpha \in (0, 1)$ and low-type with probability $1-\alpha$. At the beginning of the game, each firm chooses Q, which represents the quality of the items that it will produce throughout the game. We assume that Q takes values in the range $[\underline{q}, \overline{q}] \subseteq [0, 1]$. If a firm of type T chooses $Q = q$, they incur a one-time cost $C_T(q)$. We assume that cost functions are strictly increasing, and that high-type firms have strictly lower marginal costs of quality than low-type firms, i.e. $0 < C_H^{'}(q) < C_L^{'}(q)$.

After firms have chosen their quality, a signal S is drawn for each firm. Then, potential customers arrive sequentially and decide whether or not to purchase from the firm. If a customer buys from a firm with quality level q, he gets a reward of q and the firm gets a reward of 1. Otherwise, the customer gets a reward of $r \in (0, 1]$, which we refer to as the customer's "reserve value" or "outside option." We assume that customers know their reserve values, and it is common knowledge that these values are drawn independently and identically from a distribution with cdf F.

If the customer patronizes the firm, they leave a review of their experience, which can be either positive or negative. The probability of a positive review is equal to the quality of the item purchased, and the review is independent from all other randomness in the system. We assume that these reviews have been summarized into the sufficient statistic $X = (n_-, n_+)$, where n_-, n_+ are the number of negative and positive reviews left so far, respectively.

In this paper, we consider three information structures:

- Full observability: Customers see a signal $S = Q$ which completely reveals the firm's quality level. No additional inference can be drawn from X.
- Partial observability: Customers see the initial signal S and the reviews X.
- No observability: Customers do not see S or X when making their decision.

3.2 Equilibrium Concept

We focus on symmetric pure strategy Bayesian equilibria of this game. These are characterized by a firm strategy σ and a customer strategy ψ. The firm strategy maps type T to the quality level Q selected. Because there are only two firm types, σ is fully characterized by the values $q_H := \sigma(H)$ and $q_L := \sigma(L)$. Customer strategies map the observed elements of the triple (r, S, X) to an action $A \in \{0, 1\}$ where $A = 1$ represents a decision to purchase the item. We assume that firms and customers are risk-neutral and that customers are short-lived. Furthermore, we restrict our attention to equilibria in which the following holds:

Assumption 1. *Whenever a customer is indifferent about whether to patronize the firm, he does so. Whenever a firm is indifferent among a set of quality levels, it chooses Q to be the highest quality in this set.*

We make this assumption to simplify the notation and analysis; without it, firm best response sets may be empty. In what follows, we sometimes refer to the set of "equilibria" of a game. When doing so, we mean "symmetric pure strategy Bayesian equilibria in which Assumption 1 holds."

We define $E_\sigma[\cdot]$ (respectively, $E_\psi[\cdot]$) be the expectation of its argument when all firms play strategy σ (customers play strategy ψ). Analogously, define $P_\sigma(\mathcal{A})$ ($P_\psi(\mathcal{A})$) to be the probability of $\mathcal{A}$ if all firms play strategy σ (all customers play strategy ψ). We let $\mathcal{G}$ be the set of games described above.

3.3 Model Analysis

The majority of this paper considers separating equilibria (i.e. equilibria in which $q_H \neq q_L$) in partial information games. To set the stage for the discussion in sections 4 and 5, we establish here two results that hold for both the one-shot model and the infinite horizon model. The first states that high-type firms always choose a quality level that is at least as high as the one chosen by low-type firms. The second proposition addresses equilibria of "full observability" and "no observability" games (which are used as benchmarks in the remainder of the paper), as well as pooling equilibria of games with partial observability.

Proposition 1. *For any game $G \in \mathcal{G}$, if $\sigma = (q_L, q_H)$ is an equilibrium strategy, then $q_H \geq q_L$.*

This holds because if a low-type firm were weakly better off from playing $q_L > q_H$, then the assumption $C'_H(q) < C'_L(q)$ implies that a high-type firm strictly prefers q_L to q_H.

Proposition 2. *Let $G \in \mathcal{G}$.*
(i) If quality is unobservable, the only equilibrium is $(\underline{\sigma}, \underline{\psi})$ where $\underline{\sigma}$ satisfies $q_H = q_L = \underline{q}$ and $\underline{\psi}(S, X) = \mathbf{1}\{r \leq \underline{q}\}$.
(ii) If quality is fully observable, in equilibrium each customer purchases the item if and only if the observed signal exceeds their outside option ($r \leq S$). Therefore, firms either choose $Q = \underline{q}$ or $Q = r$.
(iii) If quality is partially observable, the pair $(\underline{\sigma}, \underline{\psi})$ described in part (i) is always an equilibrium. Furthermore, if $\overline{q} < 1$, this is the unique pooling equilibrium.

Claim (i) follows because when quality is unobservable, consumer choices are independent from the quality set by the firm. Therefore, for any fixed customer strategy, firms prefer to minimize costs by stetting the lowest quality level $\underline{q}$. The best response for customers in this case is to purchase if and only if $r \leq \underline{q}$. The second part of the proposition follows from the fact that if $\underline{q} < r$ and the firm produces at a level in the interval $(\underline{q}, r)$, it receives no customers. Similarly, there is no benefit to producing above r. Finally, (iii) follows since, when $q_H = q_L$, customers know the quality of the item they're considering. Because we have specified that indifferent customers purchase the item, this means that customer decisions don't depend on the observed review. Thus, there is no incentive for firms to choose any quality level higher than $\underline{q}$.

4 One-Period Model

4.1 Description of Equilibria

We now discuss a simple model in which at most a single customer patronizes each firm. For now, we focus on the case where costs are linear, with $C_T(q) = c_T q$, and where product quality is partially observable. This means that the customer sees S before making a decision, and that $S = 1$ with probability Q and $S = 0$ otherwise.

Some may ask where the first review S comes from. Although there are possible explanations (for example, it may be written by a professional reviewer whose job it is to try the products of new firms), this is not an essential feature of our model. Indeed, all of the intuition and techniques used below would apply to a model in which two customers consider the firm in sequence. In that model the first customer may choose not to buy from the firm, but if he does make a purchase, he leaves a review that the second customer sees. Our choice to make the first review "automatic" serves only to clarify the exposition.

We search for equilibria by fixing customer behavior ψ, computing the firm's best response, and then checking to see if this induces the specified behavior ψ. Note that ψ specifies for each r, what a customer with reserve r will do upon seeing a positive review, and what they will do upon seeing a negative review. We define $p_+(\psi)$ to be the probability that a customer playing ψ buys from the firm if $S = 1$ (i.e. the expectation of $\psi(r, S = 1)$ over possible reserves r), and $p_-(\psi)$ the corresponding probability when observing $S = 0$. Then the firm's best response to ψ is to select

$$q_T \in \arg\max_q \{qp_+(\psi) + (1-q)\,p_-(\psi) - c_T q\}\,, T \in \{L, H\}$$

Note that the objective above is linear in q, so we have a simple characterization of firm best responses. If $p_+(\psi) - p_-(\psi) - c_T < 0$, choosing $Q = \underline{q}$ is uniquely optimal. Otherwise, $\overline{q}$ is in the firm's best response set. It follows that the only separating equilibria satisfying Assumption 1 have $q_L = \underline{q}, q_H = \overline{q}$.

For any firm strategy σ, the optimal customer response is to purchase if and only if the expected quality of the item given the observed signal exceeds the value of their outside option, i.e. $E_\sigma[Q|S = s] \geq r$. In particular, if σ satisfies $q_L = \underline{q}, q_H = \overline{q}$, then

$$E_\sigma[Q|S{=}1] {=} \frac{\alpha\overline{q}^2 + (1-\alpha)\underline{q}^2}{\alpha\overline{q} + (1-\alpha)\underline{q}} \text{ and } E_\sigma[Q|S = 0] = \frac{\alpha(1-\overline{q})\overline{q} + (1-\alpha)(1-\underline{q})\underline{q}}{\alpha(1-\overline{q}) + (1-\alpha)(1-\underline{q})}.$$

The leads to the following:

Proposition 3. *If $G \in \mathcal{G}$ is a one-period game with linear costs, the only possible firm strategy in a separating equilibrium of G is $q_H = \overline{q}$ and $q_L = \underline{q}$. This outcome is supported in equilibrium if and only if $c_H \leq c_0 < c_L$, where $c_0 = F\left(\frac{\alpha\overline{q}^2+(1-\alpha)\underline{q}^2}{\alpha\overline{q}+(1-\alpha)\underline{q}}\right) - F\left(\frac{\alpha(1-\overline{q})\overline{q}+(1-\alpha)(1-\underline{q})\underline{q}}{\alpha(1-\overline{q})+(1-\alpha)(1-\underline{q})}\right)$ is a constant that depends on model primitives.*

Most of this proposition has already been established above. To complete the proof, note that if ψ is the customer best response to $(q_L, q_H) = (\underline{q}, \overline{q})$ satisfying Assumption 1, then

$$p_+(\psi) = F\left(\frac{\alpha\overline{q}^2 + (1-\alpha)\underline{q}^2}{\alpha\overline{q} + (1-\alpha)\underline{q}}\right), p_-(\psi) = F\left(\frac{\alpha(1-\overline{q})\overline{q} + (1-\alpha)(1-\underline{q})\underline{q}}{\alpha(1-\overline{q}) + (1-\alpha)(1-\underline{q})}\right).$$

Note that when c_L and c_H are both higher or both lower than c_0, no separating equilibrium exists in this model. When both firm types have high costs, attracting new customers costs more than the benefit it provides. When both firm types have low costs, all firms would like to play $\overline{q}$. If $\overline{q} < 1$, however, this cannot be an equilibrium, as discussed in Proposition 2. The problem is that in this case, customers do not get any information from the signal S; firm costs do not sufficiently differentiate firms of opposite types.

Intuitively, the farther apart c_H and c_L are, the "more likely" a separating equilibrium is to occur. We formalize this by noting the following consequence of Proposition 3.

Remark 1. If $[c_H^1, c_L^1] \subseteq [c_H^2, c_L^2]$ and $(\alpha, \underline{q}, \overline{q}, F)$ are such that a separating equilibrium exists in the partial-information game when firm costs are c_H^1 and c_L^1, then the same $(\alpha, \underline{q}, \overline{q}, F)$ also admit a separating equilibrium when firm costs are c_H^2 and c_L^2.

4.2 Welfare Comparison

Here we compare producer and consumer surplus under the different information settings. The producer surplus is taken to be the equilibrium expected profit of the average firm, while consumer surplus is the equilibrium expected utility of an average consumer. We consider the case where customers have a common reserve r, so that F is a point mass. We will later discuss how the results and intuition extend to other suitably restricted distributions F. We summarize our findings as follows:

Proposition 4. *If all customers share a common outside option $r \in \mathbb{R}$:*
(i) Both producer and consumer surplus are minimized when quality is completely unobservable.
(ii) High type firms always prefer fully observable quality, while low-type firms may prefer that quality is only partially observable.
(iii) Consumer surplus under partial observability is always at least as large as in the full information model. In any separating equilibrium, consumers are strictly better than in full information except in the case where the value of the outside option r is exactly $\frac{\alpha\overline{q}^2+(1-\alpha)\underline{q}^2}{\alpha\overline{q}+(1-\alpha)\underline{q}}$.
(iv) There are equilibria of games with partially observable quality in which both consumer and producer surplus exceed their levels in the equilibrium of a corresponding game with fully observable quality.

Claim (i) follows from Proposition 2 (i), which revealed that the unique equilibrium in the completely unobservable case is when $q_L = q_H = \underline{q}$. To see why (ii) is true, note first that since $q_L \geq q_H$ in equilibrium by by Proposition 1, it must be that $E_\sigma [Q|S=0] \leq E_\sigma [Q|S=1] \leq q_H$. That is, customers always expect to receive a lower quality good than what is truly being offered by the high-type firms. It follows by monotonicity of F that high-type firms earn no more than $F(q_H) - C_H(q_H)$ in this equilibrium. Furthermore, $F(q_H) - C_H(q_H) \leq \max_q \{F(q) - C_H(q)\}$, the amount that high-type firms earn in full information.

Now, consider claim (iii). In the full information setting, no firm will ever produce a quality level $q > r$, since by instead producing a quality strictly between r and q, it could reduce costs while still ensuring business. Therefore, customers earn exactly r in any equilibrium of the full information game.

In the partial information game, customers can still earn a certain payoff of r, and therefore expect to earn at least r in any equilibrium. In particular, given a fixed firm strategy σ, a customer with outside option r who responds optimally to σ earns $P_\sigma(S=1) \max(E_\sigma[Q|S=1], r) + P_\sigma(S=0) \max(E_\sigma[Q|S=0], r) \geq r$. Moreover, this surplus strictly exceeds r as long as $E_\sigma[Q|S=1] > r$.

As discussed in Proposition 3, the only possible separating equilibrium satisfies $q_H = \overline{q}$, $q_L = \underline{q}$. When firms play this strategy, $E_\sigma[Q|S=1] = \frac{\alpha\overline{q}^2 + (1-\alpha)\underline{q}^2}{\alpha\overline{q} + (1-\alpha)\underline{q}}$, which must be at least r in order for σ to be supported in equilibrium. Thus, unless this exactly equals r, customers are strictly better off than in the full information game.

Finally, note that although high-type firms prefer full information, low-type firms may prefer the partial information game. This is because a "lucky" signal may cause customers to mistakenly purchase from them. When $\overline{q}$ is large and c_H is sufficiently small, high-type firms are barely worse off in the partial information game, implying that firms are better off on average. For a specific numerical example take $r = 0.5,\ \underline{q} = 0.3,\ \overline{q} = 0.8,\ \alpha = 0.5,\ c_L = 6,\ c_H = 0.3$.

4.3 Discussion

To simplify the analysis we have focused on the case where firms have linear costs, and in the welfare analysis, on the case where consumers share a common outside option. The intuition and results extend beyond these restrictive cases, however.

When quality is fully observable, a firm's marginal benefit from increasing quality comes through the corresponding increase in $F(q)$. For many distributions, this benefit diminishes quickly beyond some threshold, giving firms little incentive to produce above that level. When quality is partially observable, the benefit to increasing Q comes from the rise in the probability of receiving a positive review (and the associated jump in the probability of attracting a customer). For many cost structures, and many choices of F, the differing incentives under these two information structures can lead to equilibria in which high-type firms set higher quality levels in the partial information game. We give one such example below.

Example 1. Suppose that customer reserve values are uniformly distributed on $[1/4, 3/4]$, and that $(\underline{q}, \overline{q}) = (1/4, 1)$. Suppose that firm costs are linear in quality with $c_L > 2$ and $c_H < 4/3$. Then the full information equilibrium is $q_L = 1/4, q_H = 3/4$, whereas by Proposition 3, the game with partial observability has an equilibrium in which $q_L = 1/4, q_H = 1$.

5 Infinite Horizon Model

We now consider a model in which an infinite sequence of homogeneous consumers visit each firm and firms seek to maximize expected discounted profit, with discount factor $\delta < 1$. We let X_t denote the review history seen by the t^{th} customer to consider a firm. Furthermore, we take the signal S to be uninformative, so customer inference is based only on X_t. We look for an equilibrium by fixing the firm strategy $\sigma = (q_L, q_H)$, determining the optimal customer response, and then verifying that given the induced customer behavior, no firm has an incentive to deviate from σ.

5.1 Customer Best Response

Fix a choice of σ satisfying $q_H > q_L$. The best response of a customer with reserve r is to purchase if and only if they expect that Q is at least r, given history X. Define the rejection set $R_\sigma(r)$ to be the set of review histories (n_-, n_+) such that when firms play strategy σ, the optimal decision of a customer with reserve r is not to purchase. In other words,

$$R_\sigma(r) = \{(n_-, n_+) \in \mathbb{N}^2 : E_\sigma[Q|X = (n_-, n_+)] < r\}.$$

Then the best customer response to σ is defined by $\psi(r, S, X) = \mathbf{1}(X \notin R_\sigma(r))$. It turns out that we can precisely describe these rejection sets. Note that

$$E_\sigma[Q|X] = P_\sigma(T = H|X)q_H + P_\sigma(T = L|X)q_L.$$

Since $P_\sigma(T = L|X) = 1 - P_\sigma(T = H|X)$, rearranging terms gives us that $X \notin R_\sigma(r)$ if and only if $P_\sigma(T = H|X) \geq \frac{r - q_L}{q_H - q_L}$. More algebra reveals that this is equivalent to the condition

$$\frac{P_\sigma(T = H|X)}{P_\sigma(T = L|X)} \geq \frac{r - q_L}{q_H - r}.$$

Fortunately, the ratio of conditional probabilities on the left has a nice form. By Bayes' theorem,

$$\frac{P_\sigma(T = H|X)}{P_\sigma(T = L|X)} = \frac{P(T = H)}{P(T = L)} \frac{P_\sigma(X|T = H)}{P_\sigma(X|T = L)} = \frac{\alpha}{1 - \alpha} \left(\frac{q_H}{q_L}\right)^{n_+} \left(\frac{1 - q_H}{1 - q_L}\right)^{n_-}.$$

Taking logarithms, we see that a consumer purchases if and only if

$$n_+ \log\left(\frac{q_H}{q_L}\right) - n_- \log\left(\frac{1 - q_L}{1 - q_H}\right) \geq \log\left(\frac{1 - \alpha}{\alpha} \frac{r - q_L}{q_H - r}\right). \qquad (1)$$

Thus, the optimal customer strategy takes the following appealingly simple form. Each firm starts with a "reputation score" (the left side of (1)) of zero.

Every positive review improves the firm's reputation by a fixed constant, while negative reviews decrease its reputation by a different constant. Customers map their reserve to a reputation cut-off (the right side of (1)) and purchase precisely when the firm's score is at least their cutoff. Note that the different coefficients on n_+ and n_- reflect the fact that positive and negative reviews may contain different amounts of information. If, for example, most customers have positive experiences ($q_H > q_L > 1/2$), someone reading reviews may (rationally) be said to learn more from a new negative review than a new positive one.

Put slightly differently, the firm reputation score follows a random walk. Every time the firm gets a customer, its score increases by $\log(q_H/q_L)$ with probability Q and decreases by $\log((1-q_L)/(1-q_H))$ with probability $1-Q$. The rejection region $R_\sigma(r)$ is the set of points lying below a line whose slope depends only on σ. We now point out two particularly simple instances of this decision rule.

Remark 2.
(i) If, before reading reviews, a consumer is indifferent about whether to purchase the product (i.e. $r = \alpha\, q_H + (1-\alpha)q_L$), the optimal decision rule is to go if and only if $\frac{n_+}{n_- + n_+} \geq \beta$, where β is defined to be

$$\frac{\log((1-q_L)/(1-q_H))}{\log(q_H/q_L) + \log((1-q_L)/(1-q_H))}.$$

(ii) If $q_H + q_L = 1$ and $r \in (q_L, q_H)$, the optimal customer strategy is to go if and only if $n_+ - n_- \geq d(r)$, where

$$d(r) = \log\left(\frac{\alpha}{1-\alpha}\frac{r-q_L}{q_H-r}\right) / \log\left(\frac{q_H}{q_L}\right).$$

Both of these are established via basic algebraic manipulation of (1). Note that part (i) says that an initially indifferent customer should purchase whenever the "average review" is at least β. The decision rule in part (ii), meanwhile, compares the number of positive and negative reviews in absolute terms.

5.2 Firm Revenue Calculation

Because the reserve r is identical across customers, as soon as a firm is rejected by a customer, we know that it will be rejected by each subsequent customer. Let τ be the first period in which $X_\tau \in R_\sigma(r)$, i.e. the period at which the firm has effectively gone out of business. If this never occurs, define $\tau = \infty$. The following proposition addresses a case in which it is possible to derive a closed-form expression for the firm's expected discounted profit.

Proposition 5. *Suppose that all customers go unless there are at least $m \in \mathbb{N}_+$ more negative reviews than positive ones. The firm's expected discounted profit when choosing $Q = q$ is given by*

$$\frac{1 - E[\delta^\tau | Q = q]}{1-\delta} - C_T(q), \;\; T \in \{L, H\}, \tag{2}$$

where

$$E[\,\delta^{\tau}|Q=q]=\left(\frac{1-\sqrt{1-4q(1-q)\delta^{2}}}{2\delta q}\right)^{m}. \tag{3}$$

The expression (2) arises because the firm's discounted revenue is given by $1+\delta+\cdots+\delta^{\tau-1}=\frac{1-\delta^{\tau}}{1-\delta}$. A proof of (3) appears in the Appendix.

The expression in (3) is strictly decreasing in q, suggesting that expected discounted revenue is continuously increasing in q on the entire interval $[\underline{q},\overline{q}]$. This occurs because customers only purchase from, and therefore only review, firms with sufficiently high reputation scores. As a result, even firms choosing quality levels above the customer threshold must fear that several poor reviews will leave them unable to signal their quality. This gives firms an incentive to set quality levels exceeding r, just as they had in the one-period model.

The following example illustrates that decision rules of the form assumed by Proposition 5 can be supported in equilibrium.

Example 2. When $\alpha\in(1/2,3/5), C_T(q)=c_T/(1-q), c_H=\frac{8}{9}-\frac{28}{9\sqrt{19}}\approx 0.175, c_L=\frac{9}{2}-\frac{18}{\sqrt{19}}\approx 0.371, \delta=1/2, r=1/2$, there is an equilibrium in which $q_H=0.6, q_L=0.4$, and customers purchase if and only if there are at least as many positive reviews as negative ones.

To verify this, note that given customer behavior, (2) and (3) imply that the firm best response σ is given by $q_H=0.6, q_L=0.4$. When firms choose these values, by Remark 2 (ii), customer behavior is optimal.

5.3 Welfare Analysis

Here, we discuss two welfare properties of equilibria in this infinite-horizon model. Proposition 6 says that customers who arrive later are better off. In Proposition 7 we show that customers are weakly better off in the partial information game than when quality is visible, and that some customers are strictly better of so long as high-type firms play above $\underline{q}$.

Proposition 6. *For any σ, if customers play a best response to σ, the expected surplus of the t^{th} customer to consider the firm is non-decreasing in t. Furthermore, the t^{th} customer is strictly better off than the first if the probability that the t^{th} customer buys lies in $(0,1)$.*

Proof. Let ψ be a best response to σ. The payout to a customer of type r who sees history X and plays ψ is $\max(E_{\sigma}[Q|X],r)$. Regardless of customer strategies, the sequence $\{E_{\sigma}[Q|X_t]\}_{t\geq 1}$ is a martingale. Because max is a convex function, Jensen's inequality implies that a customer of type r is weakly better off arriving in period t than in any earlier period. Averaging over r proves the first part of the proposition. Additionally, for any fixed r, if $0<P_{\psi}(r\leq E_{\sigma}[Q|X_t])<1$, then Jensen's inequality is strict, i.e.

$$E_{\psi}[\max(E_{\sigma}[Q|X_t],r)]>\max(E_{\psi}[E_{\sigma}[Q|X_t]],r)=\max(E_{\sigma}[Q],r).$$

Note that the left side represents the expected profits of customer t, and the right side equals the expected profit of the first customer.

Proposition 7. *In an infinite-horizon model where customers have identical reserves $r \in (\underline{q}, \overline{q})$, the equilibrium expected surplus of the t^{th} customer in the incomplete information game is at least as great as the equilibrium surplus of a customer in the full-information game for all t. Furthermore, unless the partial information equilibrium satisfies $q_H = q_L \in \{\underline{q}, 1\}$, this inequality is strict for sufficiently large t.*

Proof. As with Proposition 4, in the full information game firms have no incentive to set a quality above r, so all customers receive r and thus they are at least as well off in the partial information game. Since, by Proposition 2 (iii) the only possible pooling equilibria occur at $\underline{q}$ and 1, we must have $q_H > q_L$. If this is a best response, it must be that the first customer buys, which in turn implies that $E_\sigma[Q] = \alpha q_H + (1 - \alpha) q_L \geq r$. In any separating equilibrium, customers must use reviews to make decisions, i.e. there must exist review histories X that occur with positive probability such that $E_\sigma[Q|X] < r$. Thus, for sufficiently large t, the probability that the t^{th} customer buys lies strictly between zero and one. By Proposition 6, this customer is strictly better off in the partial information game than in the game of full information.

6 Conclusion

Here we have presented a simple model of reputation and product quality in markets where consumers publicly share reviews of their experience. We emphasize a setting in which customers' experiences are intrinsically random, but are positively correlated with the quality of the product. We arrive at the insight that a noisy review process of this form may yield equilibria in which firms produce higher quality goods than they would if quality were directly observable. This occurs because even high quality firms are at risk of losing future customers due to bad initial reviews. The effect is compounded by a cascade-like phenomenon: when customers are unlikely to patronize the firm, it can be difficult or even impossible for firms to improve their reputation. Above, we illustrate these ideas through stylized models. Due to the tractability of our models, we consider this paper a promising foundation for future work.

Perhaps the most natural extension of our work is to consider models in which each firm's quality level is allowed to vary over time. Models of this form can pose significant technical challenges, as strategies and the corresponding inferential procedures become much more complicated. Successful analysis of a model incorporating these elements could shed light on the question of how a firm's long run quality level depends on its review history.

Another interesting avenue for future work would be to consider a model in which firms are competing more directly with one another. In our work, firm decisions only indirectly affect other firms, by changing customer inferences. An alternative model might allow customers search, at some cost, for a firm with better reviews. Such a model has a similar flavor to the one we study, except that customer reserves are not exogenously specified, but rather determined endogenously by competition.

References

[1] Allen, F.: Reputation and product quality. RAND Journal of Economics 15(3), 311–327 (1984), http://EconPapers.repec.org/RePEc:rje:randje:v:15:y:1984:i:autumn:p:311-327
[2] Banerjee, A.V.: A simple model of herd behavior. The Quarterly Journal of Economics 107(3), 797–817 (1992), http://qje.oxfordjournals.org/content/107/3/797.abstract, doi:10.2307/2118364
[3] Bar-Isaac, H.: Reputation and survival: Learning in a dynamic signalling model. The Review of Economic Studies 70(2), 231–251 (2003), http://restud.oxfordjournals.org/content/70/2/231.abstract, doi:10.1111/1467-937X.00243
[4] Bikhchandani, S., Hirshleifer, D., Welch, I.: Learning from the behavior of others: Conformity, fads, and informational cascades. Journal of Economic Perspectives 12(3), 151–170 (1998), http://ideas.repec.org/a/aea/jecper/v12y1998i3p151-70.html
[5] Bose, S., Orosel, G., Ottaviani, M., Vesterlund, L.: Monopoly pricing in the binary herding model. Economic Theory 37(2), 203–241 (2008), http://www.jstor.org/stable/40282924 ISSN 09382259
[6] Cooper, R., Ross, T.W.: Prices, product qualities and asymmetric information: The competitive case. Review of Economic Studies 51(2), 197–207 (1984), http://ideas.repec.org/a/bla/restud/v51y1984i2p197-207.html
[7] Fudenberg, D., Maskin, E.: The folk theorem in repeated games with discounting or with incomplete information. Econometrica 54(3), 533–554 (1986), http://www.jstor.org/stable/1911307 ISSN 00129682
[8] Fudenberg, D., Levine, D., Maskin, E.: The folk theorem with imperfect public information. Econometrica 62(5), 997–1039 (1994), http://www.jstor.org/stable/2951505 ISSN 00129682
[9] Hörner, J.: Reputation and competition. American Economic Review 92(3), 644–663 (2002), http://ideas.repec.org/a/aea/aecrev/v92y2002i3p644-663.html
[10] Ifrach, B., Maglaras, C., Scarsini, M.: Monopoly pricing in the presence of social learning. Working Papers 12-01, NET Institute (August 2012), http://ideas.repec.org/p/net/wpaper/1201.html
[11] Luca, M.: Reviews, reputation, and revenue: The case of yelp. com. Com (September 16, 2011). Harvard Business School Working Paper (12-016) (2011), http://hbswk.hbs.edu/item/6833.html
[12] Rogerson, W.P.: Reputation and product quality. The Bell Journal of Economics 14(2), 508–516 (1983), http://www.jstor.org/stable/3003651 ISSN 0361915X
[13] Rubinstein, A.: Equilibrium in supergames with the overtaking criterion. Journal of Economic Theory 21(1), 1–9 (1979), http://ideas.repec.org/a/eee/jetheo/v21y1979i1p1-9.html
[14] Shapiro, C.: Consumer information, product quality, and seller reputation. The Bell Journal of Economics 13(1), 20–35 (1982), http://www.jstor.org/stable/3003427 ISSN 0361915X
[15] Smallwood, D.E., Conlisk, J.: Product quality in markets where consumers are imperfectly informed. The Quarterly Journal of Economics 93(1), 1–23 (1979), http://www.jstor.org/stable/1882595 ISSN 00335533

[16] Spence, M.: Job market signaling. The Quarterly Journal of Economics 87(3), 355–374 (1973), http://www.jstor.org/stable/1882010 ISSN 00335533

[17] Wehrum, K.: How businesses can respond to criticism on yelp (June 2009), http://www.inc.com/magazine/20090601/how-businesses-can-respond-to-criticism-on-yelp.html

[18] Wolinsky, A.: Prices as signals of product quality. The Review of Economic Studies 50(4), 647–658 (1983), http://www.jstor.org/stable/2297767 ISSN 00346527

A Computing Expected Firm Revenue

In this section, we prove Proposition 3, i.e. that when customers all choose to buy from the firm unless $n_- - n_+ \geq m \in \mathbb{N}$, and τ is defined as in Section 5,

$$E[\delta^\tau] = \left(\frac{1 - \sqrt{1 - 4q(1-q)\delta^2}}{2q\delta}\right)^m \quad \text{for } \delta \in (0,1). \tag{4}$$

We start by noting that the firm's reputation score can be re-normalized to be a simple random walk starting at 0, where τ is the first time that the walk reaches $-m$ (or ∞, if it never does). The firm chooses Q, i.e. the probability that the walk moves up.

Let $X_i \in \{-1, +1\}$ be i.i.d. and take the value +1 with probability q, $Y_n = \sum_{i=1}^n X_i$, and $\tau = \min\{n : Y_n = -m\}$. Define $\phi(\theta) = E[e^{\theta X_1}] = qe^\theta + (1-q)e^{-\theta}$, and let $\psi(\theta) = \log \phi(\theta)$. It follows that for any θ, $e^{\theta Y_n - n\psi(\theta)}$ is a martingale. Then for all n,

$$E[e^{\theta Y_{\tau\wedge n} - (\tau\wedge n)\psi(\theta)}] = 1. \tag{5}$$

Let $\theta_0 = \min\{\theta : \phi(\theta) = 1\}$. We can compute directly that if $q \leq 1/2, \theta_0 = 0$, and if $q > 1/2$ $\theta_0 = \log(\frac{1-q}{q}) < 0$. We will let $n \to \infty$ in Equation (5) and show that for $\theta < \theta_0$,

$$e^{m\theta} = E[\phi(\theta)^{-\tau}]. \tag{6}$$

Once this has been established, we use the fact that $\phi(\theta) = qe^\theta + (1-q)e^{-\theta}$ is decreasing on $(-\infty, \theta_0)$ and onto $(1, \infty)$. In particular, given $\delta \in (0,1)$ we can find a unique $\theta \in (-\infty, \theta_0)$ such that $1/\delta = \phi(\theta)$, i.e. $qe^{2\theta} - e^\theta/\delta + (1 - q) = 0$. Apply the quadratic formula to solve explicitly for θ; it is defined by $e^\theta = \frac{1-\sqrt{1-4q(1-q)\delta^2}}{2q\delta}$. Substituting $1/\delta$ for $\phi(\theta)$ in (6), this implies that $E[\delta^\tau] = \left(\frac{1-\sqrt{1-4q(1-q)\delta^2}}{2q\delta}\right)^m$, as claimed.

We now justify (6). Note that for $\theta < \theta_0$, $\psi(\theta) > 0$, so $e^{-(\tau\wedge n)\psi(\theta)} \leq 1$. Additionally, $S_{\tau\wedge n} \geq -m$. Combining these shows that $0 \leq e^{\theta Y_{\tau\wedge n} - (\tau\wedge n)\psi(\theta)} \leq e^{-m\theta}$. By the dominated convergence theorem, we may exchange limit and expectation to obtain

$$E[\lim_{n\to\infty} e^{\theta Y_{\tau\wedge n} - (\tau\wedge n)\psi(\theta)}] = 1. \tag{7}$$

If $q \leq 1/2$, $\tau < \infty$ almost surely, and thus the left side of (7) is $e^{-m\theta} E[\phi(\theta)^{-\tau}]$, which can be rearranged to yield (6).

If instead $q > 1/2$, $P(\tau = \infty) > 0$ so

$$E[\lim_{n\to\infty} e^{\theta Y_{\tau\wedge n} - (\tau\wedge n)\psi(\theta)}] = E[\lim_{n\to\infty} e^{\theta Y_n - n\psi(\theta)}; \tau = \infty] + E[e^{\theta Y_\tau - \tau\psi(\theta)}; \tau < \infty].$$

On the event $\{\tau = \infty\}$, however, $Y_n < 0$. Since $\theta < 0$, we have that $e^{\theta Y_n - n\psi(\theta)} \leq e^{-\theta m}\phi(\theta)^{-n}$. Note that since $\phi(\theta) > 1$ for $\theta < \theta_0$, $e^{-\theta m}\phi(\theta)^{-n} \to 0$ as $n \to \infty$. It follows that the left side of (7) equals $e^{-m\theta} E[\phi(\theta)^{-\tau}; \tau < \infty]$, which in turn equals $e^{-m\theta} E[\phi(\theta)^{-\tau}]$, completing the proof that (6) holds whenever $\theta < \theta_0$.