BIROn - Birkbeck Institutional Research Online

Birkbeck
UNIVERSITY OF LONDON

Customer Anger and Incentives for Quality Provision

Anthony Heyes
University of Ottawa

Sandeep Kapur
Birkbeck, University of London

August 2012

Customer Anger and Incentives for Quality Provision[*]

Anthony Heyes
University of Ottawa

Sandeep Kapur[†]
Birkbeck, University of London

August 2012

Abstract

Emotions are a significant determinant of consumer behaviour. A customer may get angry if he feels that he is being treated unfairly by his supplier and that anger may make him more likely to switch to an alternative provider. We model the strategic interaction between firms that choose quality levels and anger-prone customers who pick their supplier based on their expectations of suppliers' quality. Strategic interaction can allow for multiple equilibria including some in which no firm invests in high quality. Allowing customers to voice their anger on peer-review fora can eliminate low-quality equilibria, and may even support a unique equilibrium in which all firms choose high quality.

Keywords: anger, customer attrition, quality
JEL codes: D03, D11, L15

1 Introduction

There is much evidence that emotions – in addition to individual preferences over product attributes – are important determinants of consumer choices. The range of negative emotions that have been linked empirically to consumer behaviour include anger and disappointment (Smith and Bolton (2002)), shame and envy (Richins (1997)), and embarrassment (Ruth et al. (2002)). A range of positive emotions have also been explored – including arousal (Steencamp et al. (1996)), joy (Nyer (1997)) and gratitude (Ruth et al. (2002)).[1] As well as allowing for improved explanation of behaviour, attention has also been focused on how emotions can be manipulated by marketing and advertising (Dellarocas (2003), Bagozzi et al. (1999)).

Our focus is on anger. We develop a model of consumer behaviour in which anger emerges in response to poor service experience. We explore how anger can drive consumers to switch providers and the extent to which such customer attrition motivates firms to invest in quality to prevent poor service experience. We also consider how these features may be affected if customers share their negative experiences through online fora, in what is described as electronic-Word-of-Mouth (e-WOM).[2]

Consider a setting in which customers have incomplete information about firm characteristics, in particular firms' investment in reducing the probability of service failures. A customer – even a completely dispassionate one – who experiences a service failure may revise down his assessment of his current provider's type and switch supplier for purely 'rational' reasons. This

---

[1]Laros and Steencamp (2005) present a hierarchical meta-analysis drawing a distinction between a superordinate level of positive and negative *affect* – a general feeling of positivity or negativity – and a subordinate level of specific emotions of the sort listed above. For an excellent survey of research on the role of affect on consumer choices see the chapter by Cohen et al. (2008) in the *Handbook of Consumer Psychology*.

[2]See Bougie et al. (2003) for an early survey of the literature. Most people can think of instances where they or their acquaintances have switched their bank, mechanic, or telephone company because of anger with service deemed to be unreasonably poor.

is a *cognitive* process and corresponds to the notion of disconfirmation – the extent to which perceived quality fails to match pre-purchase expectations.

Bad service may, in addition, generate an *emotional* response. "When customers are dissatisfied, they also develop emotions, such as anger – the emotion most loaded with energy" (Chebat and Ben-Amor (2005)). Anger is a hostile emotion that is targeted at the 'cause' of the dissatisfaction, with a desire to punish the wrong-doer. As Berkowitz et al. (2004) put it in their survey, "anger is linked associatively with an urge to injure". This is consistent with the neurological evidence of de Quervain et al. (2004), amongst others, that pleasure centers in the brain are activated when those who have previously acted selfishly in an experimental economic transaction are themselves harmed. Similarly,

> "One way of thinking about this emotion in utility terms is to see angry people as people whose utility increases when the target of their anger is harmed." Rotemberg (2008: 10).

"Anger in response to a service failure arises when customers appraise an event as unfair, with high service provider control over the service failure" (Ruth et al. (2002)). Further – and important in motivating our work – anger is found to have a major impact on consumer switching. In their paper *Angry Customers Don't Come Back: The Experience and Behavioral Implications of Anger and Dissatisfaction in Services*, Bougie et al. (2003) find that "... empirically anger is a powerful predictor of customers' behavioral responses to failed service encounters (complaining and switching) *over and above the effect of dissatisfaction*". They use interview data to generate evidence on consumer emotional processes: "Angry customers wanted to 'get back at' the organization, to hurt the business of the service provider" (Bougie et al. (2003: 382)). These are what psychologists refer to as 'emotivational' goals and are commonly associated with anger in psychological studies.

Consumers react to their own experiences, but angered customer may

2

also share their negative experience by posting critical reviews on (for example) internet review sites.[3] The explosion in the use of the internet over the last two decades has made substantial changes to how the reputation of a firm for service quality evolves. Jim Lecinski – director of sales at Google – coined the term "Zero Moment of Truth" (ZMOT) to describe the moment that a consumer goes online to gather information about a good or service they are thinking of buying, noting that 70% of US consumers say that they look regularly at reviews posted on consumer websites to help with shopping (Lecinski (2008)). This presents challenges for firms and has made it important to understand how online reputations are formed through electronic word-of-mouth (e-WOM) – see, for examples, Gruen et al. (2006) and Park and Lee (2009). Empirical evidence points to the importance of e-WOM as a motive for brand choice and switching.[4]

Despite plentiful evidence of the role of anger in shaping consumer behaviour there have been relatively few attempts to incorporate it into formal economic models. Notable exceptions are papers by Rotemberg (2003, 2008 and 2011).[5]

Our basic model is presented in Section 2. Firms can position themselves in either a 'basic' segment (low price/low quality) or 'branded' segment (high price/(claimed) high quality). Consumers get angry if they feel they have been treated unfairly – if they believe their provider has charged a high price without making a good faith effort to deliver high quality. But even a firm

---

[3]There is a growing body of empirical evidence of the role of anger in motivating hostile internet reviews (see, for example, Hennig-Thurau et al. (2004) and Sparks and Browning (2010)). Shen et al. (2011) provide survey-based evidence noting that "to post is a good way to vent anger".

[4]Lecinski goes on to develop the implications of e-WOM and the ZMOT for marketing practice and how businesses should approach online reputation.

[5]Parts of our modeling strategy borrow from Rotemberg, but our interests are different. Our focus is on anger in response to lapses in service quality, and how such responses might impact the incentives for firms to invest in quality, and we analyze a competitive (rather than single firm) setting which allows us to think about churn.

that has invested in quality will occasionally suffer lapses in service standards, so a customer can infer dishonest behaviour only imperfectly. An angered customer may abandon their provider in order to punish it. We model this process of anger arousal and the subsequent switching decision quite carefully. The analysis recognizes that the desire to switch will combine rational elements with emotional ones, and that the degree of anger may depend on the price paid. We assess the extent to which the threat of consumer switching can generate incentives for provision of quality. We find that when customers learn purely from their own experience, not all premium-segment providers will invest in high quality, even when the costs of investment in quality are relatively low.

In Section 3 we introduce 'voice' – the idea that angry customers may also vent their anger to others through, say, electronic-Word-of-Mouth. The forum for such communication may vary with setting, but we have in mind websites such as Tripadvisor.com (travel) and TopTable.com (restaurants) that allow users to share reviews with others. The introduction of such social learning changes the equilibria in the model qualitatively. In particular, an outcome in which all premium-segment firms provide high-quality service – ruled out in Section 2 – can be sustained as an equilibrium. In Section 4 we outline some implications of our model and consider extensions. Section 5 concludes.

2   Model

We consider a two-period model in which a population of firms compete to offer a service to a mass of consumers. In the initial period each firm chooses a price level (high or low) for the service. In addition, some firms may be able to make a costly investment in the provision of quality. Consumers pick a provider based on the price and their expectation of quality. At the start of period 2 consumers may switch providers, and in the model that decision

4

will depend on a mixture of rational and emotional responses to their first period experience with their provider.[6]

Some firms have a fixed technology that allows them to provide only visibly low quality, q_ℓ. This segment is competitive with price p_ℓ equal to marginal cost c of providing the service, so (p_ℓ, q_ℓ) describes the 'basic' (un-branded) segment of the market.

Other firms are able to choose low or high quality provision (q_ℓ or q_h). Once chosen, quality is fixed across periods, perhaps embodied in capital, organizational structures or practices that cannot readily be adjusted. Choosing high quality q_h requires a one-off investment $\kappa_j > 0$ for firm j (this varies across firms) but marginal cost is unchanged.[7]

The branded or premium segment comprises firms that charge $p_h > p_\ell$ and claim to provide quality $q_h > q_\ell$. Customers cannot observe directly whether a particular firm has made the investment in quality, so the claim of premium status can be made dishonestly. If so, the premium segment of the market comprises some firms described by (p_h, q_h) and others described by (p_h, q_ℓ). In what follows we use $\lambda \in [0, 1]$ to denote the fraction of the former – good-faith suppliers of high quality – in the premium segment.

The price in the premium segment p_h is formed endogenously in our model, based on customer perceptions of average quality in this segment. Our model examines incentives for provision of quality by characterizing equilibrium values for λ.

---

[6]The two-period setting ensures tractability but is clearly a simplification. Later in the paper we discuss the implications of allowing for more than two periods.

[7]An alternative specification could assume that quality choice affects *marginal* (rather than fixed) cost, but this does not alter our results qualitatively.

2.1 Quality and customer experience

A firm's quality impacts its customers' experience in a 'noisy' manner – high-quality firms are more likely to deliver a good experience than low-quality firms, but do not *always* do so. Realistically, high quality does not mean perfect.

> "Mistakes are an inevitable part of every service activity. Hard as they try, even the best service companies can't prevent the occasional late flight, burned steak, or missed delivery. The fact is that in services, no matter how rigorous the procedures and employee training, or how advanced the technology, zero defects is an unattainable goal." (Hart et al. (1990: 148)).

Treating customer experience in any particular service episode as being either good (denoted as g) or bad (b), we assume a high-quality firm delivers a good experience with probability $t < 1$, while a low-quality firm delivers it with probability $s < t$. These probabilities are known to be independently and identically distributed and can be summarized as follows:

	Quality choice	
	q_h	q_ℓ
$\Pr(good|q)$	t	s
$\Pr(bad|q)$	$1 - t$	$1 - s$

The noisy relationship between firms' quality choices and customers' experience implies that customers face a signal extraction problem. A customer who experiences a bad service episode in the premium segment may have picked a low-quality firm or may just have had an unlucky encounter with a high-quality firm. Given prior beliefs that a randomly-picked provider is high quality with probability λ, the Bayesian posterior conditional on a bad experience is

$$\Pr(q_h|b) = \frac{\lambda(1 - t)}{\lambda(1 - t) + (1 - \lambda)(1 - s)} \equiv \rho_b. \tag{1}$$

Clearly $\rho_b < \lambda$ for any $0 < \lambda < 1$: a bad experience lowers confidence in the provider's quality. The gap between t and s determines the precision of the signal: with a larger gap a customer's experience is a tighter signal of the firm's quality choice.

2.2 Consumer utility

A consumer gets more utility from a good service episode than a bad one, so that $u_g > u_b$. To save notation we define $\hat{u} = u_g - u_b$. Given the relative likelihood of good and bad experiences, a consumer is willing to pay as much as $su_g + (1-s)u_b$ for service provided by a low-quality firm. We assume that customers employ at least the basic version of the service in each period.[8] With competitive prices $p_\ell = c$ in the basic segment, we only require

Assumption 1 *(Customers always purchase)* $\quad su_g + (1-s)u_b > c$.

In contrast, the expected utility of service from a reliably high-quality provider is $tu_g + (1-t)u_b$. Consumers who pick a provider randomly from the premium segment, where fraction λ of providers are expected to be high quality, are willing to pay a premium

$$p(\lambda) = \lambda(t - s)\hat{u}. \tag{2}$$

over what they pay for the service in the basic segment.

Other things equal, a consumer's choice of segment would be based on a comparison of the quality premium $p(\lambda)$ with the actual price differential $p_h - p_\ell$ across the segments. If $p_h - p_\ell < p(\lambda)$ the gain in expected utility more than compensates for the higher price p_h in the premium segment. If $p_h - p_\ell > p(\lambda)$ customers would be better-off picking a provider in the basic segment. In the absence of other considerations, both segments will co-exist

---

[8]This allows us to focus on customer churn – their movement between providers – though it could be relaxed without disturbing our results.

if and only if $p_h - p_\ell = p(\lambda)$. With the basic segment serving as a competitive fringe, with $p_\ell = c$, price p_h in the premium segment is an increasing function of λ

Of course, in two-period settings such as ours, consumers' choice in the initial period may be more complicated than a simple trade-off between current utility and price. To the extent they anticipate the possibility of switching providers in the future, they must consider their optimal *sequence* of choices. For instance, a forward-looking customer – especially one with low switching costs – might be tempted to 'take a chance' on the premium segment, knowing that in the event of a bad experience he could easily switch to the basic segment. The introduction of emotivational factors may complicate this further: consumers who know themselves particularly prone to anger might shy away from the premium segment to forestall the possibility of future anger ('anger aversion'). If so, the price differential across the segments in the initial period may depend on consumers' switching costs and on the psychic costs and benefits of their choices. While the actual price differential in the initial period $p_h - p_\ell$ may be higher or lower than $p(\lambda)$, it would nevertheless be an increasing function of λ.

The trade-off between utility and price allows us to pin down, more precisely, the price that can be sustained in the premium segment in the second period. We next focus on how that price affects customer churn within the branded segment.

2.3   Switching when consumers get angered by ripoffs

How does a customer who has paid the higher price p_h to obtain the premium service respond to an episode of bad service? As a Bayesian updater, he revises downwards his assessment that his service provider is high-quality (recall equation (1)). This generates incentives to switch providers, and we differentiate between rational incentives and emotional ones.

Rational gain from switching: Following a bad service episode, the consumer's posterior assessment that his provider is high-quality is ρ_b while the probability that a randomly-picked alternative provider is high-quality remains $\lambda > \rho_b$. So what we will call the 'rational gain from switching' is

$$(\lambda - \rho_b)(t - s)\hat{u}. \tag{3}$$

Anger and the emotivational gain from switching: Poor service may also arouse anger and emotionally-driven behaviour. Customers who feel 'ripped off' may become angry and want to get back at the firm that they believe has treated them badly. They may do this by switching – taking their custom elsewhere and so depriving the supplier of future business. A rip-off in the current setting involves a firm that charges a premium price $p_h > p_\ell$ despite being low-quality.[9] In our setting customers never know for sure that they have been ripped off, but those who have suffered a bad experience have a probabilistic assessment, $(1 - \rho_b)$. In diverting their custom away they can impose a penalty on the rip-off supplier by depriving it of future profit. For a premium-segment firm profit per customer depends on the price differential that can be sustained without losing customers to the basic segment. Assuming, for simplicity, that some consumers can switch costlessly to the basic segment, profit per customer in the second period is (at least) $p(\lambda)$. So $(1 - \rho_b)p(\lambda)$ captures the expected penalty that can be imposed upon rip-off merchants – something from which anger-prone customers derive utility or emotivational benefits (recall the earlier quotations from Rotemberg (2008) and de Quervain et al. (2004)).

---

[9]If anger is triggered by a sense of being treated unfairly, we require agents to have a reference point as to what constitutes fair treatment (this is an application of Kahneman, Knetsch and Thaler's (1986) theory that consumers feel entitled to their "reference transaction"). Firms selecting (p_h, q_h) or (p_ℓ, q_ℓ) are not only being honest but also deliver positive expected utility to the consumer. In contrast, a firm choosing (p_h, q_ℓ) is dishonest and might deliver negative utility. A customer learning ex post that their supposedly premium supplier chose low quality would feel 'ripped off' – in effect duped into making a choice by an unscrupulous firm in pursuit of excess profit.

It is realistic to suppose that customers vary according to their anger 'type' – the strength of the emotivational benefits they derive from this sort of retribution. For the i-th customer that weight is labeled by $\gamma_i \geq 0$, with distribution $G(\gamma)$ in the population. Those very prone to anger have a high value of γ, whilst a 'conventional' economic agent has $\gamma = 0$. The emotivational gain from switching after a bad service experience is

$$\gamma_i \ (1 - \rho_b)p(\lambda), \tag{4}$$

or, after substituting for $p(\lambda)$ from (2),

$$\gamma_i \ (1 - \rho_b)\lambda(t - s)\hat{u}. \tag{5}$$

A decision to switch draws on both rational and emotional elements.[10] A customer will switch in the second period if and only if the rational and emotivational returns exceed the costs of switching.

We allow heterogeneity in switching costs, so that the cost of switching from one premium segment firm to another is ω_i for the i-th customer.[11] The i-th consumer will switch following a bad experience if the sum of (3) and (5) exceeds ω_i: that is, if and only if

$$[(\lambda - \rho_b) + \gamma_i(1 - \rho_b)\lambda] \ (t - s)\hat{u} \geq \omega_i. \tag{6}$$

---

[10]Before we proceed it is worth noting an alternative interpretation of expression (4) – namely that the extent of a customer's anger depends upon the *size* of the rip-off. If a customer comes to believe that his supplier has misrepresented a low-quality offering as being high quality then the intensity of his anger might reasonably be increasing in the size of the premium. While our argument for the inclusion of $p(\lambda)$ in (4) was as a measure of amount of retribution (consistent with the desire for revenge motive highlighted in the seminal analysis of Rotemberg and others) it equally allows for this alternative interpretation. Indeed, we could model emotivational benefit as any arbitrary, increasing function of $p(\lambda)$, $(1 - \rho_b)$ and γ_i. Our results are robust to such alternative formulations.

[11]That switching costs vary across individuals is realistic, but is not strictly necessary for the analysis in this section. The assumed heterogeneity is analytically useful in later sections of this paper, where it ensures switching behaviour is always continuous in λ.

Using equation (1), this is equivalent to

$$m(\lambda) \left[(t - s) + \gamma_i(1 - s)\right](t - s)\hat{u} \geq \omega_i, \tag{7}$$

where for ease of notation we define

$$m(\lambda) \equiv \left[\frac{\lambda(1 - \lambda)}{\lambda(1 - t) + (1 - \lambda)(1 - s)}\right].$$

The term $m(\lambda)$ captures how the incentives to switch depend on λ as it varies in the unit interval. Notably, $m(\lambda)$ is non-monotonic in λ, with its value approaching zero as λ tends to 0 or to 1, but is positive in the interior of the unit interval. The result below follows.

Result 1 *Other things being equal, the gains from switching are lower when average quality in the premium segment is either very low or very high.*

The gains from switching are lowest at the two extremes, though for quite different reasons. When λ is close to zero, most firms are low quality so the expected gain from switching to a rival provider is small. On the emotivational front, while the gap $1 - \rho_b$ generates anger, the anger is tempered by the feature that the low average quality supports only a small price premium $p(\lambda)$. The emotivational gain from switching is small too.

On the other hand, when λ is close to 1 almost all firms are of high quality and the consumer's loss of confidence in his supplier in the wake of a bad service episode is less pronounced (this is a familiar feature of Bayesian updating). In other words, with a high prior a customer is more likely to interpret a bad experience as an unlucky episode with a high-quality firm, rather than indicative of low quality.

Overall, the incentive to switch varies with the anger propensity parameter γ_i and the switching cost parameter ω_i. In particular, more anger-prone customers (with high γ_i) or those with lower switching costs (low ω_i) are more likely to switch. The distribution of parameters, $G(\gamma)$ for γ_i and $\Omega(\omega)$ for ω_i, are assumed to be common knowledge, continuous, and mutually independent.

2.4 Customer attrition and aggregate churn

In the initial period, consumers choose between the basic and premium segments of the market. Those who buy from the basic segment have no incentive to switch suppliers in our simple setting, so a firm that locates itself in this segment never loses customers.

Consumers who buy initially from the premium segment cannot distinguish between high and low quality providers. For simplicity we assume that all firms in this segment are similar in size so that we can normalize their initial customer base at one unit (this is relaxed in Section 4). In this premium segment the behaviour of consumers who have a good initial experience is trivial to analyze. For them positive experience reinforces confidence in the quality of their current provider so they have no rational or emotional gains from switching – they stick with their current supplier.

Among consumers who suffer a poor initial experience, some, based on rational and emotivational considerations, may choose to switch away from their current provider. For any premium-segment firm we define the *customer attrition rate* $\sigma(\lambda)$ as the fraction of customers who abandon that firm following a bad experience in the initial period – in other words, this is the fraction of customers for whom condition (7) holds.[12] Continuity of the underlying parameter distributions ensures that σ is a smooth function of λ.[13]

To understand how $\sigma(\lambda)$ varies with λ, consider (7). Since the gain from switching is non-monotonic (Result 1), so too is the attrition rate $\sigma(\lambda)$. Indeed, given $m(\lambda)$ it is easy to check that there is no attrition at all when $\lambda = 0$ or $\lambda = 1$, and there exists a $\hat{\lambda}$ in the unit interval such that the gain from switching is increasing in the interval $(0, \hat{\lambda})$ and decreasing in the

---

[12]For given λ and γ_i the fraction of customers who switch is $z(\lambda, \gamma_i) = \Omega(m(\lambda)[(t-s) + \gamma_i(1-s)](t-s)\hat{u})$. The attrition rate is then given by $\sigma(\lambda) = \int z(\lambda, \gamma_i)dG(\gamma_i)$.

[13]To avoid trivialities associated with prohibitive switching costs, we assume $\sigma(\lambda) > 0$ for at least some values of λ.

interval $(\hat{\lambda}, 1)$.[14] We summarize this as follows.

Result 2 *If a fraction λ of the providers in the premium segment are high quality, the attrition rate for consumers who suffer a bad experience is $\sigma(\lambda)$ where*

1. *$\sigma(0) = 0$ and $\sigma(1) = 0$;*

2. *$\sigma(\lambda)$ is increasing in λ for $\lambda \in (0, \hat{\lambda})$ and decreasing in $\lambda \in (\hat{\lambda}, 1)$.*

It is straightforward to see that a more anger-prone population (in the sense of stochastic dominance of $G(\gamma)$, the distribution of the anger parameter), will display higher customer attrition rates for any λ.

The movement of consumers among premium-segment firms comes exclusively from those who have a bad service experience. Among the fraction λ of firms that are high-quality, a proportion $(1 - t)$ of customers have a bad experience in period 1; among the fraction $(1 - \lambda)$ of premium segment firms that are low-quality, a fraction $(1 - s)$ do. Adding the two, and recalling that a fraction $\sigma(\lambda)$ of customers with bad experiences switch, the aggregate churn at the end of the initial period is

$$\Psi(\lambda) = [(1 - s) - \lambda(t - s)]\sigma(\lambda). \tag{8}$$

While higher average quality reduces the incidence of bad experiences (that is, the first term is decreasing in λ), the conditional attrition rate $\sigma(\lambda)$ is non-monotonic in λ. The aggregate churn $\Psi(\lambda)$ inherits this non-monotonicity. Further, aggregate churn is zero at the extremes, $\lambda = 0$ and $\lambda = 1$.

---

[14]It is straightforward to verify that $\hat{\lambda} = \frac{(1-s) - \sqrt{(1-s)(t-s)}}{t-s}$. There are two solutions to the quadratic equation that leads to this, but the other solution lies outside the unit interval so is not relevant.

2.5 Quality choice by firms

For firms that present themselves in the premium segment, quality choice involves a trade-off: high quality requires upfront investment κ_j but enhances customer retention. To the extent that retained customers are profitable this provides an incentive to invest in quality. Recall that given λ, profit per customer in the second period equals $p_h - c = \lambda(t - s)\hat{u}$.

More generally, investment in quality can improve a firm's profitability both through better retention of its existing customer base and through improved acquisition of new customers. Analysis in this section is simplified, however, by the feature that customers learn only from their own experiences and – garnering no information on rival providers – switching customers can do no better than selecting a new provider at random. This means a firm's choice of quality does not affect the rate at which it picks-up new disaffected customers at the start of period 2. We relax this assumption in Section 3.

We consider how a premium-segment firm's choice affects its two-period profits, ignoring discounting. A firm that chooses q_h expects to lose a fraction $(1 - t)\sigma(\lambda)$ of its initial customers in period 2, while a firm that chooses q_ℓ while pretending to be high quality expects to lose a fraction $(1 - s)\sigma(\lambda)$. Thus, investing in high quality reduces customer attrition by $(t - s)\sigma(\lambda)$, thereby improving future profits by amount $\lambda\sigma(\lambda)(t - s)^2\hat{u}$. At the same time choosing high quality entails higher cost κ_j so that the net incremental gain from choosing high quality is

$$\Delta(\lambda, \kappa_j) = \lambda\sigma(\lambda)(t - s)^2\hat{u} - \kappa_j. \tag{9}$$

Crucially, a firm's incremental profit depends on λ (as the aggregate proportion of high-quality providers affects its customer retention) and on κ_j (its direct cost of investing in high quality).

A firm will prefer to choose high quality, then, if and only if $\Delta(\lambda, \kappa_j) \geq 0$. For any λ, this condition is more likely to be met for firms whose cost of investing in quality is relatively low, that is, those with κ_j small enough. Let

14

the distribution of κ_j be given by a continuous distribution function $K(\kappa)$. Then, for any λ, the fraction of firms for whom investment in high quality is profitable is given by $F(\lambda) = K(\lambda\sigma(\lambda)(t-s)^2\hat{u})$.

Note that firms' incentives to choose high quality depend on expected attrition, but customers' attrition rates vary themselves with average quality λ in the premium segment (recall Result 2). Given this strategic interaction, we focus on Nash equilibria in which (a) each firm's choice of quality is optimal given the choices of other firms and behaviour of customers and, (b) customers switch optimally given their beliefs about firms' quality. If $F(\lambda)$ captures the mass of firms that find if profitable to choose high quality when aggregate proportion of high quality firms is believed to be λ, equilibrium configurations require that $F(\lambda^*) = \lambda^*$.

Our setting allows for multiple equilibria. We begin with one which is degenerate in the provision of quality. Consider a configuration in which no firm chooses high quality, that is, consider $\lambda = 0$. From Result 2, there is no customer attrition at all and if investment in quality is costly, it follows from equation (9) that it will not be profitable for any firm to invest in quality. We have an equilibrium with $F(0) = 0$.

If the fixed costs of investing in quality is relatively high for all firms, the above outcome will be the unique equilibrium: if $F(\lambda) < \lambda$ for all $\lambda > 0$, attrition rates are too low to justify any investment in high quality to be sustained in equilibrium. To explore outcomes other than this trivial one, we consider the case where costs of investment in quality are low enough and anger-fueled churn sufficiently strong, so that $F(\lambda) > \lambda$ for at least some λ. Formally, we have

Assumption 2 *There exists some $\lambda \in (0, 1)$ such that $F(\lambda) > \lambda$.*

This assumption allows for the existence of multiple equilibria.

Proposition 1 *In a setting where angry customers may switch to a randomly-picked rival provider, strategic interaction between customers and firms gen-*

15

erates multiple equilibria. The equilibria differ in the quality choice of firms in the premium segment and associated customers' attrition rates.

- *There exists an equilibrium in which all firms choose low quality ($\lambda = 0$) and the attrition rate for disgruntled customers is zero ($\sigma(0) = 0$).*

- *Under Assumption 2 there also exist equilibria in which some fraction $\lambda^* \in (0,1)$ of firms choose high quality and disgruntled customers' behaviour is described by $\sigma(\lambda^*) > 0$.*


Figure 1: Fraction of firms choosing high quality

A formal proof is in the Appendix. But to understand the qualitative properties of these equilibria consider Figure 1, which plots $F(\lambda)$, the fraction of firms that find it profitable to invest in quality, as a function of λ. Given the pattern of customer attrition described in Result 2, $F(\lambda)$ is zero at $\lambda = 0$ (and, indeed, with strictly positive κ_j, for all λ sufficiently close to zero) and also at $\lambda = 1$. For some intervening range in the unit interval, investment in quality is profitable for low-cost firms, so that $F(\lambda)$ is strictly positive. Assumption 2 places the stronger requirement that $F(\lambda) > \lambda$ for some λ: if so, the graph of $F(\lambda)$ must intersect the 45-degree line.

Equilibria are characterized by λ^* such that $F(\lambda^*) = \lambda^*$. At the first equilibrium – labeled e_0 in the figure – all firms choose low quality ($\lambda_0^* = 0$). If so, customers do not gain from switching (both the rational and emotivational gain from switching are zero). Without customer attrition, there is no incentive for any firm to invest in quality, so that $F(\lambda_0^*) = 0$. This equilibrium is degenerate from the viewpoint of quality provision and characterized by *no-quality-and-no-churn*. Effectively, in this case the premium and basic segments of the market coincide.

When Assumption 2 holds, the interaction between firms and angry customers supports other (possibly multiple) equilibria, here labeled e_1 and e_2. At these equilibria customer attrition is such that a fraction $F(\lambda^*) = \lambda^*$ of firms – those whose fixed costs of investment in quality are low enough – find it profitable to invest in high quality. As shown here, this equality obtains at two distinct values of λ, with $\lambda_2^* > \lambda_1^* > 0$. However, the equilibrium associated with the lower of these values is unstable, leaving e_2 as the locally stable equilibrium. Importantly, this equilibrium is characterized by positive churn and positive investment in quality.

The role of customer anger in this setting merits clarification. A more anger-prone population of customers – in the sense of stochastic dominance of $G(\gamma)$ – will lead to higher conditional attrition rates $\sigma(\lambda)$ and thereby increase the number of firms that find it profitable to invest in high quality for any λ: in terms of the equilibria described above, $F(\lambda)$ is weakly increasing in anger. If so, an angrier population is associated with higher equilibrium value λ_2^*. Put simply, greater customer anger sustains greater provision of quality.

Besides, anger may have a discrete effect on the equilibrium outcomes if it implies a greater likelihood that Assumption 2 holds: in other words, there may be parameter configurations where attrition rates for dispassionate customers are too weak to support *any* investment in quality (that is, e_0 is the unique equilibrium) while angry customers might well induce some

investment in quality (support e_2 as an equilibrium).

However, remarkably, anger cannot support comprehensive investment in quality, even when the population is extremely prone to anger and the cost of investment in quality is relatively small. Formally,

Result 3 *An outcome in which all firms in the premium segment choose high quality is never an equilibrium.*

To see why, note that if all firms choose high quality, there is little incentive for customers to switch providers. But if the attrition rate is zero, firms' incentive to invest in high quality disappears. Thus all premium-segment firms choosing high quality cannot be an equilibrium outcome in this setting.

3 Revenge of keyboard warriors

Our analysis in the previous section shows that when anger manifests itself in the form of 'exit' alone, it provides only limited incentives to invest in quality. Equilibrium outcomes may involve no investment in quality or, at best, even when the cost of investing in quality is relatively low, less than complete investment in quality.

In many settings disgruntled customers react to poor service not just by switching their own custom, but also by sharing their negative experiences with others. Historically this may have been through negative word-of-mouth within a limited community, but the internet has allowed for much wider dissemination. Online sharing of bad experiences has become quite common. Some of the most popular websites are essentially fora for customers to share experiences with strangers (for example, tripadvisor.com), but social networking tools are often used to share information about errant service providers. As the BBC's Business Editor put it:

18

"Once upon a time companies could afford to be rude. Angry customers would grumble to a few friends, withdraw their custom, but there was little else they could do. Today, they still tell their friends, but they do it online, using media websites like YouTube, Facebook and Twitter"

BBC Business News website, 3 October 2010.

Here we extend our model to allow for public sharing of negative experiences.[15] Significantly, we find that it impacts not only the intensity of customer attrition, but also the qualitative characteristics of the equilibria. The outcome depends upon the nature and extent of learning in ways that we will describe.

In terms of our model, incorporating an assumption that some (or all) angry customers disseminate negative reviews has two effects: (a) It allows customers to learn more about the quality-type of their *own* provider. Rather than relying just on their own experience, customers can learn from the negative reports of others. Was my negative experience at hotel X an isolated instance, or did other guests also have bad experiences? Second, (b), it also allows customers to learn more about the quality-type of *other* providers. This means that switchers, rather than going to a randomly-chosen alternative, can make a more informed choice of new supplier. In particular they can identify firms that are more likely to be high quality. So switchers are disproportionately 'channeled' to high-quality firms. From a firm's perspective, investment in quality boosts profits both through better retention of current customers and better acquisition of new ones.

To assess the impact of these information channels, we develop two cases. We begin, most starkly, with the case of complete information in period 2, before considering a more general case.

---

[15]We ignore dishonest reviews. Firms may be tempted to use dishonest reviews to boost their own standing, or to sabotage the standing of competitors (Dellarocas (2006)).

3.1 Complete information

Assume that all consumers who have a bad service experience in the initial period post a negative review on a public forum. With a sufficiently large number of customers, such reviews would be completely revealing. Suppose at the end of the initial period every consumer knows the quality type of every firm, including their own, with certainty. This is a stark and unrealistic assumption, but it will provide a benchmark.

What does switching behaviour look like in this case? At the end of the first period, customers who find themselves with low-quality providers have the usual rational and emotivational incentives to switch. The rational gain from switching is now $(t - s)\hat{u}$ (because a switcher is able to identify a new provider who is high-quality with probability one). The emotivational benefits are now $\gamma_i(t - s)\hat{u}$. The switching condition, then, is

$$(1 + \gamma_i)(t - s)\hat{u} \geq \omega_i. \tag{10}$$

Comparing (10) with (6), it is easy to see that the customers are more likely to switch in this case. The associated attrition rate, which we denote as $\overline{\sigma}$, is higher than the previously-derived $\sigma(\lambda)$ for all λ.

In this setting investment in quality is attractive to a forward-looking firm both because it reduces customer attrition and also enhances acquisition of new customers who have deserted other providers. With complete information a low-quality firm expects to lose a fraction $\overline{\sigma}$ of its customer while a high quality firm loses none. The advantage of higher quality in terms of customer retention is precisely $\overline{\sigma}$.

The gain from acquiring new customers depends on the aggregate churn. Given attrition rate $\overline{\sigma}$ for low-quality firms and zero for high quality ones, aggregate churn in a population where a fraction λ is high quality equals

$$\overline{\Psi}(\lambda) = (1 - \lambda)\overline{\sigma}. \tag{11}$$

We assume that switchers are allocated equally among firms that are revealed to be of high quality. Combining the gains from better (in fact, complete)

customer retention and better customer acquisition, high quality improves period-2 customer base by

$$\overline{\sigma} + \frac{\overline{\Psi}(\lambda)}{\lambda} = \frac{\overline{\sigma}}{\lambda} \ . \tag{12}$$

Given that firms' quality levels are known in period 2, high-quality firms sell at a premium $p(1) = (t - s)\hat{u}$, with implied profit $(t - s)\hat{u}$ per customer. The incremental net benefit from investing in high quality equals

$$\overline{\Delta}(\lambda, \kappa_j) = \frac{\overline{\sigma}}{\lambda}(t - s)\hat{u} - \kappa_j. \tag{13}$$

Comparing (9) and (13) it is easy to see that $\overline{\Delta}(\lambda, \kappa_j) > \Delta(\lambda, \kappa_j)$: complete information boosts both customer retention and acquisition, generating stronger incentives for firms to invest in high quality. Since the attrition rate $\overline{\sigma}$ is positive, $\lim_{\lambda \to 0} \overline{\Delta}(\lambda, \kappa_j)$ is positive for any (finite) κ_j.

As before, given the distribution of fixed costs κ_j, we can compute the fraction of firms that will find it profitable to invest in high quality. We define this as $\overline{F}(\lambda) = K(\frac{\overline{\sigma}}{\lambda}(t - s)\hat{u})$. Observe – in contrast to the case in Section 2 – that $\overline{F}(\lambda)$ is *weakly decreasing* in λ, with $\lim_{\lambda \to 0} \overline{F}(\lambda)$ strictly positive. We assume κ_j are bounded above so, in fact, $\lim_{\lambda \to 0} \overline{F}(\lambda) = 1$.

We can note immediately – and in contrast to Proposition 1 – that the outcome in which no firm invests in quality is not an equilibrium. Intuitively, if no other firm invests in high quality, it is extremely profitable for any particular firm to do so – with full information it would expect to attract all switching customers. In this case the equilibrium is necessarily unique.

Proposition 2 *With complete information, the strategic interaction between customers and firms generates a unique equilibrium: it involves greater customer churn and a higher fraction of firms choosing high quality than the case without complete information.*

Once again, a formal proof is in the Appendix. Importantly, here the firms' choice of quality is constrained only by the cost κ_j. We cannot rule

21

out the possibility that if this fixed cost is low enough for all firms, they would *all* find it profitable to invest in quality. In that case, equilibrium is given by $\overline{F}(1) = 1$.

3.2 Incomplete information: a general case

The outcome of social learning in Section 3.1 is stark – at the end of period 1 all consumers have full information. In realistic settings the extent of learning from the experiences of others is less comprehensive, and we model this here in a very general extension. We consider a setting in which each customer learns about his provider's quality both from his own private experience and through other channels of public information, such as online reviews. The precise informational content of these public channels may vary, and different individuals may see, assimilate and combine information from different sources to varying extents. We develop an argument directly in terms of attrition rates, so that the details of the learning mechanism are 'black-boxed'.

Assume that customers assess their current provider's quality both directly (through their own experience) and through social learning (that collates others' experiences). For a firm that choose quality $q \in \{q_\ell, q_h\}$ we define $\sigma_{q\epsilon}(\lambda)$ to be the average attrition rate for customers whose own direct experience is $\epsilon \in \{b, g\}$. Specifically, for a low-quality firm, the attrition rate is $\sigma_{\ell b}(\lambda)$ for customers who have a bad experience, and $\sigma_{\ell g}(\lambda)$ for those who had a good experience. Similarly, for a firm that chooses high quality, we can define $\sigma_{hb}(\lambda)$ and $\sigma_{hg}(\lambda)$ as the attrition rates contingent on bad and good experiences. As before these attrition rates will vary with λ.

We can compare these attrition rates quite generally. With the possibility of learning from others, even customers with a positive personal experience with their provider may sometimes be tempted to switch to other providers who have garnered even more favorable reviews: in other words, $\sigma_{qg}(\lambda)$ could well be positive in this setting. But other things equal, a firm faces higher

attrition from customers who have a bad experience rather than a good one, as the former are likely be more pessimistic and angrier: we have $\sigma_{\ell b}(\lambda) \geq \sigma_{\ell g}(\lambda)$ and $\sigma_{hb}(\lambda) \geq \sigma_{hg}(\lambda)$.

To the extent customers also learn from others' experiences, they are more likely to be pessimistic about – and hence more likely to abandon – firms whose low quality is revealed through others' reviews. We expect $\sigma_{\ell b}(\lambda) \geq \sigma_{hb}(\lambda)$ and $\sigma_{\ell g}(\lambda) \geq \sigma_{hg}(\lambda)$. In words, for any category of direct experience, the availability of public information is likely to reinforce attrition from low-quality relative to high-quality providers.[16]

Given the differential attrition rates we can evaluate aggregate churn based on the proportion of firms that choose high or low quality and the customers' experiences contingent on those quality choices. We expect

$$\Psi(\lambda) = \lambda[t\sigma_{hg} + (1 - t)\sigma_{hb}] + (1 - \lambda)[s\sigma_{\ell g} + (1 - s)\sigma_{\ell b}], \qquad (14)$$

where the conditional attrition rates $\sigma_{q\epsilon}$ themselves vary with λ.

As before, a firm's quality choice affects its future customer base through its impact on retention and on acquisition. We consider these in turn. Choosing high quality rather than low improves customer retention by

$$[s\sigma_{\ell g} + (1 - s)\sigma_{\ell b}] - [t\sigma_{hg} + (1 - t)\sigma_{hb}]. \qquad (15)$$

The acquisition rate for new customers varies with the availability of public information. In the absence of any public information about firms' quality, as in Section 2, there was no differential in the acquisition rate for a high quality firm relative to a low quality firm. When public information was completely revealing, as in Section 3.1, all switching customers were

---

[16]The scenarios analyzed earlier in the paper are nested as special cases. In particular when private experience is the only channel for learning (Section 2) we have $\sigma_{\ell b} = \sigma_{hb} = \sigma_b$ and $\sigma_{\ell g} = \sigma_{hg} = \sigma_g$; indeed, there $\sigma_b = \sigma$ and $\sigma_g = 0$. With complete information, individual experiences do not matter, so that $\sigma_{\ell b} = \sigma_{\ell g} = \sigma_\ell$ and $\sigma_{hb} = \sigma_{hg} = \sigma_h$. In the complete information version in Section 3.1, the setting was even sharper, with $\sigma_\ell = \overline{\sigma}$ and $\sigma_h = 0$.

able to choose a high quality firm with probability 1. We generalize these environments by supposing that, in a population where a fraction λ of firms is high quality, a switching customer can identify a high-quality firm with probability λ^α, where $0 \leq \alpha \leq 1$ is a measure of noise. When $\alpha = 0$, the customer is able to pick a high-quality firm with probability 1; when $\alpha = 1$, a customer can pick a high-quality firm with probability λ (that is, no better than a random pick). Importantly, the attrition rates $\sigma_{q\epsilon}$ defined above will themselves depend on α.

So, in general, a fraction λ^α of the churn $\Psi(\lambda)$ is directed to high-quality firms, while the residual fraction $(1 - \lambda^\alpha)$ ends up with low-quality firms. Given the proportions of firms in each group, the differential advantage, in terms of customer acquisition, for a high quality firm is[17]

$$\left(\frac{\lambda^\alpha}{\lambda} - \frac{1 - \lambda^\alpha}{1 - \lambda} \right) \Psi(\lambda) = \left(\frac{\lambda^{\alpha-1} - 1}{1 - \lambda} \right) \Psi(\lambda). \tag{16}$$

Combining the gain from better customer retention (15) with that from better acquisition (16), the choice of high relative to low quality improves a firm's second-period customer base by

$$\beta(\lambda) \equiv \lambda^{\alpha-1}[s\sigma_{\ell g} + (1 - s)\sigma_{\ell b}] + \left(\frac{\lambda^\alpha - 1}{1 - \lambda} \right) [t\sigma_{hg} + (1 - t)\sigma_{hb}]. \tag{17}$$

Note that $\beta(\lambda)$, which summarizes the impact of customer churn on a firm's future customer base in any particular information setting, is sensitive to customers' propensity to anger, especially when anger has a marked effect on attrition rates σ_{hb} and $\sigma_{\ell b}$ in the event of unhappy experiences for customers.

The final step is to assess the impact of the customer base on profits. If switching customers end up with high-quality firms with probability λ^α, the mark-up of premium-segment over basic-segment price that the market

---

[17]To see this, note that the churn $\lambda^\alpha \Psi(\lambda)$ is shared across the fraction λ of firms that is high quality, while the residual $(1 - \lambda^\alpha)\Psi(\lambda)$ is shared among the $1 - \lambda$ low quality firms.

will sustain in the second period is $\lambda^\alpha(t-s)\hat{u}$.[18] For firm j, the incremental benefit from investing in high quality is

$$\Delta_\alpha(\lambda, \kappa_j) = \beta(\lambda)\lambda^\alpha(t-s)\hat{u} - \kappa_j, \tag{18}$$

so that the fraction of firms who invest in high quality equals

$$F_\alpha(\lambda) = K(\beta(\lambda)\lambda^\alpha(t-s)\hat{u}). \tag{19}$$

Proposition 3 *With partial public information, strategic interaction between firms' quality choices and customer attrition allows multiple equilibria.*

- *If $F_\alpha(0) = 0$, then there exists an equilibrium in which no firm invests in high quality and there is no attrition.*

- *If $F_\alpha(\lambda^*) = \lambda^*$ for some $\lambda^* \in (0, 1]$, there exist equilibria in which a fraction λ^* of firms choose high quality and attrition rates are positive.*

This is a general result and follows directly from the previous arguments. It subsumes the previous cases but also covers a much wider set of environments.

For instance, suppose that publicly-available information from customer review sites is such that customers learn nothing beyond their personal experiences about their own providers, but if they choose to abandon their current provider the available information can guide them reliably to high-quality alternatives. Formally, we have $\alpha = 0$. Choice of high quality then enhances the future customer base by

$$\beta(\lambda) = \frac{1}{\lambda}[s\sigma_{\ell g} + (1-s)\sigma_{\ell b}], \tag{20}$$

---

[18]Again the earlier versions are nested as special cases. In particular, if $\alpha = 1$ switching customers pick another firm randomly, so will pay a premium of no more than $\lambda(t-s)\hat{u}$ (the personal experience model in Section 2). If $\alpha = 0$, they pick an unambiguously better firm and will pay a premium of up to $(t-s)\hat{u}$ (the full-information version in Section 3.1).

with incremental profit

$$\Delta_0(\lambda, \kappa_j) = \frac{1}{\lambda}[s\sigma_{\ell g} + (1 - s)\sigma_{\ell b}](t - s)\hat{u} - \kappa_j. \tag{21}$$

With public information, attrition rates are necessarily positive even for small λ, so at least some firms have incentive to invest in high quality. If so $\lim_{\lambda \to 0} F_0(\lambda) > 0$ then $\lambda = 0$ cannot be an equilibrium outcome here. At the other extreme, for λ close to 1, customers' gain from switching are limited, and with low attrition rates the incentive to invest in quality is less than complete: we have $F_0(1) < 1$. If so, $\lambda = 1$ cannot be an equilibrium either. In this case, equilibrium outcome(s) necessarily involve only some firms choosing high quality.

The second part of Proposition 3 highlights the finding that in circumstances where public information, say that triggered by anger, guides customer churn, the provision of high quality by all firms may be sustained as a possible equilibrium. Indeed, the possibility exists even when information revelation is not complete, and is in stark contrast to the result in Proposition 1 – developed in the absence of sharing of information amongst consumers – where this could never be the case.

4 Extensions

Our model is stylized. In this section we explore the implications of relaxing its assumptions for a broader assessment of the relationship between customer anger, switching behaviour, and quality choice.

4.1 Repetition

The two-period structure of our model ensures tractability but is limited in scope. Without providing a full-fledged multi-period model we conjecture

that extending the number of periods would complicate the analysis in a number of ways.

With multiple periods we expect consumers' learning and their responses to be more complex. The subjective assessment of a firm's quality would be based not on a single experience with its service but on the entire history of experiences. Consumer responses would incorporate the possibility of ongoing learning. Following an initial bad experience, a customer may choose to give his provider the 'benefit of the doubt'. Repeated bad experience may trigger anger and the search for a better provider.

Similarly the punishment options available to angered customers would be richer, as in other repeated games involving punishment. Anger might result in temporary defection ('will not patronize your restaurant for n periods') or permanent defection. We would also have to incorporate the possibility that customers' anger gradually dissipates as the emotional trigger becomes more distant in time, but equally 'forgiveness' may be an aspect of 'rational play' in the strategic interaction. If firms can adjust their quality over time, we might expect a multitude of equilibria, supported by the standard folk-theorem arguments. If, on the other hand, quality is fixed across periods, low-quality providers might be driven out progressively through time.

We could also consider the possible entry of new providers: this would throw up questions about how a new entrant wins custom and builds a reputation against incumbent providers.

Multiple periods also raise the possibility of a sequence of short-lived generations of customers interacting with long-lived firms. If the transmission of firms' reputations across generations is less than complete, periodic replenishment of the customer pool with arrival of 'new blood' may help preserve the customer base of low-quality firms. This would dilute incentives to invest in quality, in a variant on the 'tourists and natives' model of Salop and Stiglitz (1977). However, if the inter-generational sharing of experience is high, it would reinforce incentives to invest in quality.

Multi-period versions of our extended model – that with social learning – might generate some interesting additional dynamics. Our working assumption was that angry customers provide reviews based on their own experience alone. Evidence on review behaviour is more subtle. Dellarocas and Narayan (2010), for example, study some dynamic aspects of internet postings in the context of movie reviews and find that, other things equal, the propensity to make a negative post is increasing in the number of existing negative posts, suggesting the possibility of path dependence in social learning.

4.2 Other instruments to manage reputation

In our model firms choose price and quality but are otherwise passive in the face of customer anger. In real-world setting firms might be expected to manage their reputations by responding to complaints in other ways. Most consumer-facing firms have specialist departments that respond to individual instances of poor customer experience.

A variety of instruments may be used to manage angry customers – for instance, simple apologies, money-back guarantees, or discounts on future purchases. The efficacy of some of these strategies have been studied before in managerial settings.[19] Traditionally the focus of management practice in this area has been retention of the angry customer himself but preventing hostile internet postings could be a more important consideration.

One strategy to assuage anger is to offer a 'money-back guarantee' that offers a rebate m as compensation to any customer who has suffered a verifiably bad experience. If such a guarantee can prevent customer attrition altogether, firms face a simple one-period pricing problem, but one where the price is contingent on consumer experience. The expected cost of such

---

[19]See, for examples, Smith and Bolton (2002) and Menon and Dube (2004). Kubo et al. (2012), for example, present experimental evidence on how a simple apology for a bad episode reduces the physiological manifestations of anger (skin temperature, skin conductance response, etc.).

guarantees is lower for high-quality firms, so we can construct equilibrium outcomes in which only firms that can invest in quality at relatively low fixed costs κ_j do so.

Consider a setting in which firms that choose high quality charge a price p_{hg} to the fraction t of its customers who have a good experience, while charging a lower effective price $p_{hb} = p_{hg} - m$ to those who have a bad experience. Firms that choose low quality charge a fixed price p_ℓ, which in a competitive environment will be driven down to marginal cost c. Here the contingent price-schedule (p_{hg}, p_{hb}) should be such that customers are willing to buy from high-quality firms, and at the same time firms that do not invest in quality do not find it profitable to mimic the price schedule of high-quality firms. It is easy to show that $p_{hg} = c + (1-s)\hat{u}$ and $p_{hb} = c - s\hat{u}$ can support such an equilibrium. Note that such equilibria require a bad experience to be verifiable, otherwise all customers would gain from claiming a bad experience. Where money-back guarantees can support such equilibria, the provision of quality does not rely on the mechanisms described in this paper.

An alternative strategy, closer in spirit to our model of customer attrition, might involve giving disgruntled customers a discount voucher worth ϵ against their future purchase. The award of such vouchers as compensation for poor customer experience is commonplace in many markets, such as air travel. A voucher is of value only to customers who stay with the firm, so it is natural to analyze in our two-period set-up. Higher-valued vouchers can better placate angry customers so the attrition rate is a decreasing function of ϵ: we write $\sigma(\epsilon)$ with $\sigma' < 0$. Also, the prospect of future discounts ϵ in the event of unsatisfactory experiences would allow firms that offer such discounts to charge a higher price ex-ante, so that we have an increasing price function $p_h(\epsilon)$.

Our focus here is not on the optimal choice of ϵ, but on the implications of this remedy for the initial choice of quality.[20] For a firm that picks low

---

[20]A firm seeking to maximize returns from retention of dissatisfied customers must pick

quality while pretending to be high quality, second-period profits are

$$s(p_h(\epsilon) - c) + (1 - s)(1 - \sigma(\epsilon))(p_h(\epsilon) - c - \epsilon). \tag{22}$$

Here the first term represents future profits from satisfied customers and the second term reflects the profit from unhappy customers that the firm manages to retain. In comparison the choice of high quality entails higher initial cost κ_j but delivers higher future profit

$$t(p_h(\epsilon) - c) + (1 - t)(1 - \sigma(\epsilon))(p_h(\epsilon) - c - \epsilon). \tag{23}$$

The net incremental gain from investing in quality is

$$\Delta_\epsilon = (t - s)[(1 - \sigma(\epsilon))\epsilon) + (p_h(\epsilon) - c)\sigma(\epsilon)] - \kappa_j \tag{24}$$

The last expression is easy to understand. The choice of high quality reduces the number of unhappy customers by fraction $(t - s)$ which results in savings on aggregate discounts paid for active retention (the first term in square brackets) and avoids loss of profits from customers who choose to leave regardless (the second term). Recall that in our basic model, which ignored active retention, the incremental gain is simply

$$\Delta_{\epsilon=0} = (t - s)[(p_h(0) - c)\sigma(0)] - \kappa_j \tag{25}$$

Given any distribution of costs κ_j, in which setting would more firms be inclined to invest in quality ex-ante? We are unable to order the two cases unambiguously. In general the ability to placate unhappy customers ex-post reduces the incentive to invest in quality ex-ante, but as $\sigma(0) > \sigma(\epsilon)$, the comparison is sensitive to the cost of placating customers. Consider an extreme case in which optimally-chosen discount successfully eliminates all attrition. With $\sigma(\epsilon) = 0$, the choice of quality involves a tradeoff between the incremental cost of future discounts $(t - s)\epsilon$ and κ_j.

ϵ to maximize $(1 - \sigma(\epsilon))(p_h(\epsilon) - c - \epsilon)$. In our setting the optimal discount is independent of the choice of quality, so in what follows we assume that it is set at this optimal value.

In many cases future discounts are offered not to all unhappy customers but only to those who complain vociferously. Such targeting amounts to a form of price discrimination: those who reveal themselves more anger-prone (of high γ type) by complaining are offered future discounts particular to them, just as those known to have high elasticities of demand are offered lower prices by a price-discriminating firm. If anger can be targeted in this way, it further reduces the incentive to invest in quality for all.

Some review fora offer suppliers a right-of-reply in response to a hostile review (see, for example, Tripadvisor.com). As Goldman (2011) points out, this feature can allow a provider not only to correct misinformation but also to turn a simple negative statement (complaint) into a conversation between the firm and the customer. Note again that the objective is not only (or even primarily) to assuage *that* aggrieved customer, but the numerous third parties who are 'over-hearing' that conversation online. The importance of customer-to-customer transmission of grievances is evident in the effort that many firms now put into managing their online reputations. Sometimes this has proved controversial, in particular the use of software packages such as Klout and Debatescape which companies use to trawl social networking sites looking for people posting negative comments, who they then contact. As British Telecom's Director of Customer Services Warren Buckley put it recently, online chatting is like "... someone having a conversation in a pub – just a very big pub. We can't stop people saying negative things about us. What we can do is identify them and offer to address their concerns."[21]

---

[21]British Telecom uses Debatescape to find customers making negative comments online, who are then contacted by "BT Sarah". Many find this practice unnerving (see the headline in the Daily Mail, 6 June 2010, "How BT Sarah Spies on Your Facebook Account") but companies defend it as a legitimate instrument of customer relations.

4.3 Social learning and the size-distribution of firms

The impact of increased internet usage and the significance of e-WOM for consumer behaviour may have significant impact on the evolution of firms and their size distribution. There is an emerging interest and literature on this relationship.[22]

Cheung et al. (2010) survey the impact of e-WOM on sales growth of firms. Online reputations and review histories have been shown to be significant determinants of sales of electronic games (Zhu and Zhang (2010)), and box-office receipts (Duan et al. 2008), among other things. Lee et al. (2011) note a debate amongst researchers regarding the likely impact of the emergence of e-WOM on the distribution of firm size.

To explore these sorts of effects satisfactorily would require a more sophisticated specification of the 'architecture' of review websites than embedded in the model presented here. Popular review websites collate and present their information in a variety of different ways. Some websites summarize reviewers' data to rank rival service-providers – e.g. Tripadvisor.com tells the viewer that, for instance, Pappasitto's Cantina is the 14th best restaurant out of 1413 in Houston. Expedia.com gives a 'star rating' on a scale of 1 to 5 to hotels whilst Toptable.com provides a score out of 10 to restaurant experiences, so these are effectively normalized to provide a measure of complaint-intensity. Most – including those just mentioned – allow the user to access the written reviews in raw form.

The way in which consumers combine information from multiple websites with their own priors and experiences is, of course, open to question. Even in the simplified environment faced by the consumer in our base model, the updating of beliefs was a statistically sophisticated and computationally demanding process. In reality people are more likely to apply heuristics in combining different sources of evidence.

---

[22]Dellarocas (2006) provides a theoretical model of the impact of e-WOM on various market outcomes, including sales, consumer surplus and firm profitability.

While several studies have focused on the average (mean or median) review score, Clemons et al. (2006) point to the likely significance of extreme reviews. Using data from beer review websites (beerhunter.com and ratebeer.com) they show that sales growth or loss were better predicted by the numbers of extreme ratings: "... increasingly reviews are prepared and posted by individuals who have been either appalled by, or delighted by an individual product or service experience." Recognizing that posts come from individuals in the tails of the distribution of willingness to complain complicates inference, but we would still expect a Bayesian to update to prefer firms with the lowest measure of complaints-intensity, which is what is needed to drive the qualitative results in this paper.

In many cases social learning may be biased towards the reporting of *negative* reviews, in that online fora are more likely to attract angry customers eager to vent their anger than those who are delighted.[23] Further, websites such as ripoffreport.com and complaintsboard.com seek to collate consumer complaints: by design information disseminated on these websites is drawn from just one tail of the distribution. In such settings social learning has the potential to damage a firm's carefully built brand or reputation.

Indeed large firms may be more vulnerable to reputational damage through e-WOM: they are more likely to feature on freely-accessible review websites, and their larger customer base may generate greater instances of complaints. Complaints against small firms may be relatively sparse, and often confined to specialist fora such as Angie's List – limited to paying subscribers. While users of online fora should correctly normalize the number of complaints against the size of the customer base (large firms are often at pains to project the number of complaints against them in the context of their large customer base), the heuristics commonly used to assess firm reputations may not capture this fully. If so, growing importance of social learning has the potential

---

[23]Consider, for instance, the suggestive title of Blackshaw's best-seller, *Satisfied Customers Tell Three Friends, Angry Customers tell 3000: Running a Business in Today's Consumer-Driven World*, Crown, 2008.

to drive customers away from larger firms towards small firms. This would, other things equal, have the effect of reducing heterogeneity of firm size. In sum, with only negative reviews, the inability of consumers to normalize the number of complaints by the size of firm (number of transactions) could be a source of dis-economies of scale.

Our analysis so far has suggested that social learning has a positive impact on average quality. However, simple modifications of our model, while preserving its essential mechanics, can suggest alternative possibilities, including settings where the impact of social learning on average quality may be perverse.

To develop this we consider, as in our basic model, a setting in which firms pick their quality levels and customers move among firms reacting to their direct experience and information gathered through social learning. However we now extend the setting to allow for multiple periods. Typically, in multi-period settings progressive gathering and sharing of information would induce customers to converge to high-quality firms. We restrict the possibility of such complete convergence by assuming that there are overlapping generations of customers, each living for only two periods and choosing from long-lived firms. If information on firm reputations is not transmitted fully across generations, low-quality firms can persist in the long run.

We also allow for size heterogeneity among firms. In particular, we consider an industry that has N small firms each of unit size (as measured by their initial customer base) and one large firm of size $B > 1$. For any generation, consumers pick a firm based on price and prior perception of quality in the initial period. In the next period customers may switch based on their past experience and information gathered from others of their generation. Those who have a positive experience stick with their current provider. Of those who have negative experiences, a fraction switch based on rational and emotivational reasons.

As before, let σ denote the attrition rate among those who have a poor

experience with their provider. But we now allow the possibility that attrition rates may vary across firms. Specifically, assuming that social learning makes the reputations of large firms more fragile than small ones, we posit that $\sigma_L > \sigma_S$. Over time, as social learning becomes more significant as an information channel, this differential in attrition rates might grow.

Our purpose is to identify outcomes in which social learning *lowers* average quality amongst providers. We focus on an equilibrium in which the large firm finds it profitable to invest in high quality, and among the cluster of N small firms only a fraction λ do.[24] The first feature requires that for the large firm the cost of investment in quality is sufficiently low relative to the advantage it produces in terms of retention and acquisition of customers. The second feature requires that when the N otherwise-identical small firms are ordered by the cost of investment in quality, that investment is just optimal for the marginal firm. We assume that the distribution of costs supports this equilibrium outcome.

The differential in average quality between the large firm and small firms can be sustained in equilibrium only if there is a compensating price differential. If only a fraction λ of small firms are high quality, buying from a randomly-selected small firm will deliver a good experience with a lower probability $(\lambda t + (1 - \lambda)s)$ as compared to probability t when buying from a large firm. If so, small firms can sell only if they charge less than the large firm: we have $p_S(\lambda) < p_L$.

Consider how consumers might switch across firms over time. Among those who initially picked the large firm and had a poor experience, a fraction σ_L will abandon that firm, moving instead to one of the small firms. Given the large firm's size B, the aggregate mass of customers switching from the large to small firms is

$$\Psi_{LS} = B(1 - t)\sigma_L. \tag{26}$$

---

[24]This is a simple modification of our previous setting in which, with firms of identical size, a subset – the so-called 'basic segment' – were unambiguously low quality while in the premium segment a fraction λ chose high quality.

Consumers who are disappointed by their initial pick of a small firm may switch either to another small firm or to the more expensive but ostensibly high-quality large firm – recall that the price differential across the two segments leaves consumers indifferent. To simplify the analysis assume that customers abandoning small firms always move to the large firm, though the assumption is not essential for our qualitative argument.[25] In this setting, the aggregate churn from the N small firms to the large firm is

$$\Psi_{SL} = N[\lambda(1-t) + (1-\lambda)(1-s)]\sigma_S. \tag{27}$$

Over time the relative shares of the large and small firms in this industry will depend on the volume of these flows across segments. If, for instance, churn from the large firm exceeds that from small firms (that is, $\Psi_{LS} > \Psi_{SL}$) the large firm will shrink and small firms will grow in size.

In steady state, the relative shares will be stable only when $\Psi_{LS} = \Psi_{SL}$, or that

$$B(1-t)\sigma_L = N[\lambda(1-t) + (1-\lambda)(1-s)]\sigma_S. \tag{28}$$

Consider how this steady state might be altered if increased reliance on social learning caused σ_L to rise faster than σ_S. To match the increase in the left hand side of the above equality, the right hand side must rise. This may be achieved in various ways, either an increase in N, the number of firms in the small-firm segment, relative to B, or through a fall in average quality λ in the small-firm segment (recall that $(1-s) > (1-t)$, so that a fall in λ will cause the right-hand side to increase). Regardless, average quality in the industry is hurt by increased reliance on social learning amongst customers of the large firm. In plain terms, if social learning increases attrition away from the reliably high-quality large firm towards firms that are less-reliably high quality, average quality falls.

_____

[25]The assumption is easily justified in settings where small firms have limited geographical reach, so that any customer has access to only one small firm while the large firm can serve all customers.

5 Conclusions

It is widely accepted that emotions – especially anger – can be important determinants of consumer behaviour. Our ambition in this paper has been to embed anger into a formal model of supplier-switching in a repeat purchase setting, and to explore the relationship between consumer anger and the firms' provision of quality.

Our model builds upon the small literature on anger and market interaction initiated by Julio Rotemberg. Our analysis differs in several fundamental ways, however. In terms of modeling anger we adopt a continuous measure – capturing the notion of degrees of anger, consistent with psychological research – rather than the binary angry/not angry approach adopted by Rotemberg. Moreover by modeling a competitive market, rather than a single firm, we are able to think about the strategic inter-dependence among sellers. In our model, for example, the extent to which bad performance by one firm will lead to lost customers is sensitive to choices made by its competitors, and vice versa.

We find that the strategic interaction between providers and consumers can result in multiple equilibria, with the equilibria varying in the fraction of firms supplying high quality and the level of customer churn. Consumer anger, by augmenting churn, reinforces incentives for service providers to invest in the provision of quality. However, when consumers' switching is guided purely by their own experience with the service, it supports only limited investment in quality: in our basic model, not all firms invest in high quality even when investment is not very costly, and there may be equilibrium outcomes in which there is no investment in quality at all.

In contrast, social learning – say, the public sharing of negative service experiences through online media – increases attrition and also directs switching customers to firms that are more likely to be high quality. The combined effect is to induce more investment in quality: in our model the low quality

equilibrium is eliminated and, with enough social learning, we may have outcomes where all firms that are able to choose to invest in provision of high quality. While learn through social media is likely to reinforce average quality, we also outline settings where it might have perverse effects on average quality.

Our model is simple and we have outlined some of the ways in which it can be extended. Future work should allow for more nuanced management of reputation by firms using non-price signals, such as investing in a brand. But the essential elements of our model – that customers can get angry when they feel badly treated by their provider and that firms will recognize that possibility when making quality decisions – are readily defensible.

References

1. Bagozzi, R., Gopinath, M., Nyer, P., 1999. The role of emotions in marketing. *Journal of the Academy of Marketing Science* 27(2), 184-206.

2. Berkowitz, L., Harmon-Jones, E., 2004. Toward an understanding of the determinants of anger. *Emotion* 4, 107-130.

3. Bougie, R., Pieters, R., Zeelenberg, M.N., 2003. Angry customers don't come back, they get back: the experience and behavioral implications of anger and dissatisfaction in services. *Journal of the Academy of Marketing Science* 31(4), 377-93.

4. Chebat, J.-C., Ben-Amor, K., 2005. Switching costs' emotional effects on dissatisfied customers' behavior. *Capco Journal of Financial Transformation*, 21, 113-120.

5. Cheung, C.M., Thadani, D., 2010. The state of electronic word-of-mouth research: a literature analysis. *Pacis 2010 Proceedings* 151, 1580-87.

6. Clemons, E.K., Gao, G. and Hitt, L., 2006. When online reviews meet hyperdifferentiation: a study of the craft beer industry. *Journal of Management Information Systems* 23(2), 149-71.

7. Cohen, J.B., Pham, M., Andrade, E.B., 2008. The nature and role of affect in consumer behavior. In: Haugtvedt, C., Kerr, P., and Kardes, F., (Eds.), *Handbook of Consumer Psychology*. Erlbaum.

8. de Quervain, D.J., Fishbacher, U., Treyer, V., Schellmanner, M. Schnyder, U., Buck, A, Fehr, E., 2004. The neural basis of altruistic punishment. *Science* 305, 1254-58.

9. Dellarocas, C., 2003. The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science* 49(10), 1407-1424.

10. Dellarocas, C., 2006. Strategic manipulation of internet opinion forums: Implications for consumers and firms. *Management Science* 52(10), 1577-1593.

11. Dellarocas, C., Narayan, R., 2010. What motivates consumers to review a product online? Mimeo, R.H. Smith School of Business, University of Maryland, College Park.

12. Duan, W., Gu, B., Whinston, A.B., 2008. The dynamics of online word-of-mouth and product sales: An empirical investigation of the movie industry. *Journal of Retail* 84(2), 233-42.

13. Fleder, D., Hosanagar, K., 2009. Blockbuster culture's next rise and fall: The impact of recommender systems on sales diversity. *Management Science* 55(5), 697-712.

14. Goldman, E., 2011. The regulation of reputational information. In Szoka, B., Marcus, A. (Eds.) The Next Digital Decade: Essays on the Future of the Internet, TechnoFreedom: Washington DC.

15. Gruen, T.W., Osmonbekov, T., Czaplewski, A.J., 2006. eWOM: the impact of customer-to-customer online exchange on customer value and loyalty. *Journal of Business Research* 59(4), 449-56.

16. Hart, C.W., Heskett, J., Sasser Jr, W.E., 1990. The profitable art of service recovery. *Harvard Business Review* 68, 148-56.

17. Hennig-Thurau, T., Gwinner, K., Walsh, G., Gremler, D.D., 2004. Electronic word-of-mouth via consumer opinion platforms: what motivates consumers to articulate themselves on the internet? *Journal of Interactive Marketing* 18(1), 38-52.

18. Kahneman, D., Knetsch, J., Thaler, R., 1986. Fairness as a constraint on profit. *American Economic Review* 76, 728-41.

19. Klemperer, P.D., 1987. Markets with consumer switching costs. *Quarterly Journal of Economics* 102(2), 375-94.

20. Kubo, K., Okanoya, K., Kawai, N., 2012. Apology isn't good enough: An apology suppresses an approach motivation but not the physiological and psychological anger. PLoS ONE 7(3): 1-6.

21. Laros, F., Steencamp, J., 2005. Emotions in consumer behavior: a hierarchical approach. *Journal of Business Research* 58(5): 1437-1445.

22. Lee, J., Lee, J.N., Shin, H. 2011. The long tail or the short tail: Category-specific impact of eWOM on sales distributions. *Decision Support Systems* 51(4), 466-79.

23. Lecinski, J., 2008. ZMOT: Winning the Zero Moment of Truth, e-book, Google.

24. Menon, K., Dube, L., 2004. Service provider responses to anxious and angry customers: different challenges, different payoffs, *Journal of Retailing* 80(3), 229-37.

25. Nyer, P.U., 1997. A study of the relationships between cognitive appraisals and consumption emotions. *Journal of the Academy of Marketing Science* 25(4), 296-304.

26. Park, C., Lee, T.M., 2009. Information direction, website reputation and eWOM effect: A moderating role of product type. *Journal of Business Research* 62(1), 61-67.

27. Richins, M.L. (1997). Measuring emotions in the consumption experience. *Journal of Consumer Research* 24(2), 127-46.

28. Rotemberg, J.J., 2003. The benevolence of the baker: fair pricing under the threat of customer anger. Mimeo, Harvard Business School.

29. Rotemberg, J.J., 2008. Behavioral aspects of price setting, and their policy implications. NBER Working Paper 13754, Boston MA.

30. Rotemberg, J.J., 2011. Fair pricing. *Journal of the European Economic Association* 9(5), 952-981.

31. Ruth, J.A., Brunel, F., Otnes, C.C., 2002. Linking thoughts to feelings: consumption emotions in a mixed-emotions context. *Journal of the Academy of Marketing Science* 30(1): 44-58.

32. Salop, S., Stiglitz, J., 1977. Bargains and ripoffs: A model of monopolistically competitive price dispersion. *Review of Economic Studies* 44(3), 493-510.

33. Shen, W., Cai, J., Li, L., 2011. Electronic word-of-mouth in China: motivational analysis. International Conference on E-Business and E-Government (ICEE): Shanghai, China.

34. Smith, A., Bolton, R., 2002. The effects of customers' emotional responses to service failures on their recovery effort and satisfaction judgements. *Journal of the Academy of Marketing Science* 30(1), 5-23.

35. Steencamp, J.B., Baumgartner, H., van der Wulp, E., 1996. The relationship among arousal potential, stimulus evaluation and the moderating role of need for stimulation. *International Journal of Research in Marketing* 13(3), 319-29.

36. Sparks, B., Browning, V., 2010. Complaining in cyberspace: The motives and forms of hotel guests' complaints online. *Journal of Hospitality Marketing and Management* 19, 797-818.

37. Zhu, F., Zhang, X., 2010. Impact of online customer reviews on sales: The moderating role of product and consumer characteristics. *Journal of Marketing* 74(2), 133-48.

Appendix

Proof of Proposition 1: The first equilibrium is straightforward to establish. If no firm invests in quality ($\lambda = 0$), customers derive no benefit from switching, so $\sigma = 0$ is an optimal response. At the same time, if no customer switches in the face of a bad experience, the choice of low quality choice does not hurt a firm's profit and avoids cost κ_j: if so, $\lambda = 0$ is optimal.

Under Assumption 2, there exist further equilibria. To characterize these we consider the magnitude $F(\lambda) - \lambda$. From Result 2, $\sigma(1) = 0$, so $\Delta(1, \kappa_j) = -\kappa_j$: if investment in quality is not profitable when $\lambda = 1$, we have $F(1) = 0$, so that $F(\lambda) - \lambda$ is negative as λ tends to 1. By Assumption 2, $F(\lambda) - \lambda$ is strictly positive for some $\lambda \in (0, 1)$. As $F(\lambda) - \lambda$ is a continuous function of λ, with a positive value for some λ and negative value as λ tends to 1, there must exist some value λ^* such that $F(\lambda^*) - \lambda^* = 0$.

We cannot rule out multiple equilibria of this kind. Given that $\sigma(0) = 0$, then $F(\lambda) - \lambda < 0$ for λ close to zero whenever $\kappa_j > 0$. If so, there is a second equilibrium: however this equilibrium is unstable. To see why, note that for $\lambda > \lambda_1^*$ in the neighborhood of λ_1^*, we have $F(\lambda) > F(\lambda_1^*)$, which will cause λ to rise. $\qquad \square$

Proof of Proposition 2: With complete information, a firm's incremental gain from investing in quality

$$\overline{\Delta}(\lambda, \kappa_j) = \frac{\overline{\sigma}}{\lambda}(t - s)\hat{u} - \kappa_j$$

is decreasing in λ. If so, $\overline{F}(\lambda)$ the fraction that find it profitable to invest is (weakly) decreasing in λ, and hence $\overline{F}(\lambda) - \lambda$ is strictly decreasing. We see that $\overline{F}(\lambda) - \lambda$ is positive at $\lambda = 0$ (because $\lim_{\lambda \to 0} \overline{F}(\lambda) = 1$) and non-positive at $\lambda = 1$, so there must exist some λ_c^* such that $\overline{F}(\lambda_c^*) = \lambda_c^*$: this characterizes the unique equilibrium for this complete information case.

Under complete information, investment in high quality improves both customer retention and acquisition: the higher profitability of investing in quality under complete information implies $\overline{F}(\lambda) > F(\lambda)$ for all λ. If so, equilibrium under complete information supports strictly higher provision of quality.

If the fixed costs κ_j is low enough such that $(\overline{\sigma}(t-s)\hat{u}) - \kappa_j > 0$ for all i, we have $\overline{F}(1) = K(\overline{\sigma}(t-s)\hat{u}) = 1$. In that case it is a dominant strategy for every firm to invest in quality. $\qquad\square$