ISSN 1350-6722

DISCUSSION PAPERS IN ECONOMICS

Inattentive consumers and product quality

by

Mark Armstrong
Yongmin Chen

Discussion Paper 07-03

UCL Department of Economics
University College London Gower Street London WC1E 6BT

Inattentive Consumers and Product Quality*

Mark Armstrong† Yongmin Chen‡

March 2008

Abstract

This paper presents a model in which some consumers shop on the basis of price alone, without attention to product quality. A firm may "cheat" (i.e., cut quality) to exploit these inattentive consumers. In the unique symmetric equilibrium, firms follow a mixed strategy involving both price and quality dispersion. Firms are less likely to cheat when there are fewer inattentive consumers, which improves welfare, but the impact on profit and consumer surplus is ambiguous. With many sellers, approximately half of them cheat, and introducing more sellers boosts consumers surplus and reduces profit, while the impact on welfare is ambiguous.

1 Introduction

Consumers may not possess all relevant information when purchasing a product from a particular seller. The missing information could be the price of the product sold by rival sellers, the quality of the product, or the availability of alternative sellers. Sometimes a consumer's lack of information could be her rational choice in an environment of costly information gathering and processing. But at other times, a consumer's inattention to information may be due to other factors. For instance, she may be unaware of such information, or she may mis-perceive the game being played by sellers.

Varian (1980) considers a market in which some consumers do not know market prices and simply purchase from a random seller, while other consumers know the prices of all sellers and buy at the lowest price.[1] (Varian's model is presented in more detail

*We are grateful to V. Bhaskar, Benedikt Koehler, Priya Sinha and seminar participants at Ohio State University, University of California-Riverside and University of Texas-Austin for useful comments. Armstrong is grateful to the Economic and Social Research Council (UK) for funding assistance.

†Department of Economics, University College London, Gower Street, London WC1E 6BT, UK. *Email*: mark.armstrong@ucl.ac.uk.

‡Department of Economics, University of Colorado at Boulder, Boulder, CO 80309, USA. *Email*: yongmin.chen@colorado.edu.

[1] See Salop & Stiglitz (1977) and Rosenthal (1980) for related models.

in Appendix A.) In this paper, we study a model in the spirit of Varian's model in which all consumers know all prices, but only some consumers pay attention to product quality. We interpret "quality" quite broadly to include financial items such as non-salient charges, price hikes after "introductory offers", and so on. In many markets the headline price is transparent and prominently displayed, while quality can be more opaque. For instance, marketing material for a new credit card may emphasize a headline interest rate, but puts relevant information about penalty charges or when interest rates are levied in the small print which some consumers overlook. A mortgage seller may offer a generous deal for the first two years of the contract, but then move to a less-than-generous interest rate for the remainder of the term.[2] A seller of insurance may advertise a headline premium, while details about excesses and exclusions are more hidden. There are several non-financial applications of our framework too. For instance, a snack could be made expensively using good ingredients or cheaply by using lots of salt, say, but only a fraction of consumers think to look at the list of ingredients. And only a fraction of consumers have a good idea about the quality of individual wines.

Chapter 4 of Cruickshank (2000) describes a survey on consumer attentiveness in financial retail markets. For instance, Table 4.4 reports that 18% know the interest rate for unauthorized overdrafts on their own current account "exactly", while 43% know this figure "not at all". Table 4.5 reports that 29% of respondents know the notice period on their savings account "exactly" while 23% know it "not at all". (Similar figures apply to knowledge about charges for withdrawals.) That chapter also documents significant price dispersion in these markets. Other illustrative information is found in Financial Services Authority (2006, pages 99-100), which reports that when buying a financial product about half of survey respondents read the terms and conditions "in detail", while about 10% did not read them at all. Korobkin (2003, footnote 45) describes a court document reporting that AT&T found that only 30% of its customers would read its entire letter updating its contract terms, 10% would not read it all, and 25% would throw away its mailing without even opening it.[3]

[2]For instance, see the opening remarks made on March 22, 2007 by Senator Chris Dodd (Chairman of the Senate Committee on Banking, Housing and Urban Affairs) at a hearing on "Mortgage Market Turmoil: Causes and Consequences". The statement referred to "Amy Womble [...] who was promised a mortgage at $927 per month and ended up with a mortgage costing her $2,100." See *http://dodd.senate.gov*.

[3]Section II of Korobkin (2003) provides an excellent summary of psychological theory and evidence about consumer decision making. An important insight is that many consumers are selective in which

Section 4.1 of DellaVigna (2007) suggests a theoretical framework in which the true total price of an item is $V = v + i$, where v is the "visible" price and i is a less visible component (such as "postage and packing" for online retailers or indirect taxes not factored into the displayed price). Due to limited attention to i, though, consumers behave as though the total price were instead $\hat{V} = v + (1-\theta)i$, where θ is an "inattention" parameter. DellaVigna summarizes a series of empirical studies in terms of their implied estimates for θ. These include the estimated inattention parameter being in the range $0.18 \leq \theta \leq 0.45$ for sales of CDs on eBay, and in the range $0.75 \leq \theta \leq 0.94$ (extremely high figures) for indirect US state taxes on beer and groceries.

Our model is a stylized representation of these kinds of situations. In particular, we suppose there are just two levels of quality which can be chosen by a seller, where the low-quality product is "useless" and would not be produced if there were no inattentive consumers in the market. The reason a firm finds it profitable to offer a low-quality product is that some consumers mistake it for the high-quality product. The model has a unique symmetric equilibrium, in which both prices and qualities are chosen stochastically by sellers. In contrast to Varian's model, here the inattentive consumers can be genuinely "ripped off" since they may receive a strictly negative payoff from participating in the market.

Different forms of public intervention could have different effects on market parameters: *competition policy*, through forbidding mergers and encouraging entry, may increase the number of rivals in the market, while *market transparency* reforms (e.g., improving information flows to consumers, or imposing disclosure requirements which force more consumers to read the small print in contracts) could increase the fraction of attentive consumers. In Varian's model, greater market transparency always benefits both groups of consumers, but is bad for profit, while competition policy has no impact on industry profit, helps the attentive consumers, but harms inattentive consumers. (See the appendix for details.) In our model, market transparency reforms improve overall welfare (which is unaffected by policy in Varian's model) and the surplus enjoyed by the attentive consumers. However, the impact on industry profits, on inattentive consumers and on aggregate consumer surplus is ambiguous. In situations where most consumers are attentive, however, the comparative statics parallel Varian's model: profits fall and

product attributes (e.g., price but not some aspects of quality) they consider before purchase.

the inattentive consumers are better off when the fraction of attentive consumers is increased still further. Tighter competition policy, at least when there are a reasonably large number of sellers present already, acts to reduce industry profits and improve aggregate consumer surplus. Total welfare can go up or down with more sellers. Attentive consumers generally are better off with more sellers, while inattentive consumers can be made worse off (as in Varian's model). Strikingly, when there are many sellers, approximately half of them will "cheat" and offer a low-quality product, regardless of other model parameters.

In our model, the inattentive consumers are unsophisticated and make no attempt to infer product quality from a supplier's price. Indeed, in our model a firm's price completely reveals its choice of quality in equilibrium, and so Bayesian consumers could infer quality even if they cannot observe it directly. There is a rich literature investigating how prices can signal quality to Bayesian consumers who cannot observe quality directly. (See section 4.2 below for a brief discussion.) This literature is especially relevant in markets in which consumers buy repeatedly, or where for some other reason consumers have a good understanding of how the market operates.[4] For instance, most people have learnt that the price of a bottle of wine is correlated with its quality even if they know little about the quality of any particular wine producer. Our model is not appropriate for such situations. Rather, our model is intended to apply to situations where some consumers have little experience of how the market operates. Examples are "one-off" purchase decisions, including the choice of bank account, mortgage contract, insurance contract (where a supplier's obligations in the event of a claim are revealed infrequently), a home improvement contract, and so on.

Gabaix & Laibson (2006) present a related model to ours, where some (non-Bayesian) consumers do not consider the price of a complementary "add-on" product (such as in-room phone charges in a hotel) when they buy the main product (the hotel room). A major difference between their model and ours is that they assume consumers need not consume the add-on, or can find a close substitute for it. As a result, sophisticated consumers who investigate (or foresee) the high price of the add-on can avoid being exploited by high add-on prices. In our model, by contrast, the product is indivisible and consumption is a one-off decision. Gabaix & Laibson assume that suppliers cannot

[4] However, with the "salty snack" example mentioned earlier, the credence good nature of the product means that consumers may not learn about ingredients even if they buy repeatedly.

distinguish between attentive and inattentive consumers, and so must charge the same price for the main item to the two groups. Since inattentive consumers generate profits from their consumption of add-ons, suppliers in a competitive market will subsidize the main item's price in order to attract these profitable consumers. As a result, the sophisticated consumers (who avoid the exploitative add-on prices) obtain the service at a lower price than they would if all consumers were sophisticated. By contrast, in our model the attentive consumers are always harmed by the presence of inattentive consumers in the market.[5]

In Section 2 we formulate our model and characterize the market equilibrium, while the impact of two policy interventions—to increase market transparency and to increase the number of suppliers in the market—is described in Section 3. Section 4 concludes the paper by discussing two extensions of the main model: to allow for heterogenous tastes for quality, and to allow those consumers who do not directly observe quality to form sophisticated Bayesian beliefs about quality.

2 A Model with Inattention to Quality

Consider a variant of Varian (1980) where, instead of being inattentive to price, some consumers do not pay attention to potential differences in the product's quality. Suppose $N \geq 2$ identical and risk-neutral firms supply a product which is homogeneous except for the fact that it may be offered with two quality levels, high (denoted H) or low (L).[6] All consumers have the same preferences for quality. Each consumer's valuation of the low-quality product is assumed not to exceed its marginal cost; for convenience we set them equal and normalize them both to zero. The marginal cost of the high-quality product is $c > 0$, while each consumer's valuation for the high-quality product is $v > c$. Each consumer has a unit demand, and the consumer population size is normalized to 1. Each firm can choose just one quality level.

Suppose all consumers observe the product's price but only a fraction λ also observe each firm's quality. The remaining $1 - \lambda$ do not observe any firm's quality and we make

[5] See Armstrong (2008) for a discussion of when the presence of inattentive consumers harms or benefits the attentive consumers, as well as of market transparency policies in general.

[6] The binary nature of quality is not unduly restrictive in our framework. Once a firm decides to cheat on quality, i.e., produce a product which the attentive consumers will not buy, it might as well choose the lowest feasible quality since this reduces its cost of supplying the inattentive consumers.

the behavioral assumption that these inattentive consumers buy on the basis of price alone (as long as the price is no higher than v). In effect, they (mistakenly) hold the belief that unobserved quality is always H and purchase from the seller with the lowest price.

When consumers are homogeneous in terms of their information, i.e., $\lambda = 0$ or $\lambda = 1$, the unique equilibrium involves marginal-cost pricing and zero profits. When $\lambda = 1$, in equilibrium all firms supply quality H and the price is c. When $\lambda = 0$ all firms supply quality L and the price is zero.

From now on assume $0 < \lambda < 1$. Then one can show that with two or three firms in the market there is no pure strategy equilibrium (either symmetric or asymmetric). However, with more firms there is an asymmetric pure strategy equilibrium:

Lemma 1 *If $N \geq 4$, the game has no pure strategy symmetric equilibrium, but it has a pure strategy asymmetric equilibrium of the following form: at least two firms offer $\{p = c$, quality $= H\}$ and at least two firms offer $\{p = 0$, quality $= L\}$. Firms make zero profit in equilibrium and total welfare, which is equal to aggregate consumer surplus, is*

$$W = \lambda(v - c) .$$

Proof. At any potential pure strategy symmetric equilibrium, either all firms choose high quality, or all firms choose low quality. If all choose high quality, Bertrand competition implies that all firms charge $p = c$. But then a firm can benefit from deviating to low quality and a price slightly below c. If all choose low quality, Bertrand competition implies that all firms charge $p = 0$. But then a firm can benefit from deviating to high quality and $p = v$. Therefore the game has no pure strategy symmetric equilibrium.

On the other hand, if at least two firms offer $\{p = c$, quality $= H\}$ and at least two firms offer $\{p = 0$, quality $= L\}$, then all firms earn zero profit and no firm has a profitable deviation. Thus the proposed set of strategies is indeed an equilibrium. The expression for total welfare follows, since only an attentive consumer will receive positive surplus, equal to $v - c$. ■

For any $N \geq 2$, the game has a symmetric mixed strategy equilibrium, which generates positive profits. (We will see in section 3.2 that when the number of firms is large, this mixed strategy equilibrium converges to the asymmetric pure strategy outcome in

Lemma 1.) Here, each firm chooses its price p randomly on an interval $[p_0, v]$. There is a threshold price $p_1 \in [p_0, v]$ such that a firm chooses low quality for sure when it chooses price $p \in [p_0, p_1)$ and chooses high quality for sure when $p \in (p_1, v]$. (At the threshold p_1 a firm is indifferent between the two quality levels.) Denote by P the probability that any given firm "cheats", that is it sets $p \in [p_0, p_1]$ and offers the low-quality product. We will show that this cheating probability is the unique solution for P in the interval $[0, 1]$ of the equation[7]

$$\left(\frac{P}{1-P}\right)^{N-1} = \frac{1-\lambda}{\lambda}\frac{c}{v-c}\left[1+\frac{1-\lambda}{\lambda}(1-P)^{N-1}\right]. \tag{1}$$

The boundary prices p_0 and p_1 are then given in terms of P by

$$p_1 = c + \frac{\lambda P^{N-1}}{\lambda + (1-\lambda)(1-P)^{N-1}}(v-c) \leq v\ , \tag{2}$$

and

$$p_0 = p_1(1-P)^{N-1} \ \leq p_1\ . \tag{3}$$

The details of the equilibrium are described in our main result:

Theorem 1 *Suppose there are $N \geq 2$ suppliers. Let P, p_1 and p_0 be defined in (1)–(3), and define $F(p)$ to be the function such that when $p_0 \leq p \leq p_1$*

$$F(p) = 1 - \left(\frac{p_0}{p}\right)^{\frac{1}{N-1}}\ , \tag{4}$$

and when $p_1 \leq p \leq v$ it is given implicitly by

$$(1-\lambda)\left[1-F(p)\right]^{N-1} + \lambda\left[P+1-F(p)\right]^{N-1} = \lambda P^{N-1}\frac{v-c}{p-c}\ . \tag{5}$$

Then $F(p)$ is a continuous cumulative distribution function with support $[p_0, v]$, where $P = F(p_1)$. The game has a unique symmetric mixed-strategy equilibrium where firms choose prices randomly with support $[p_0, v]$ and cdf $F(p)$. A firm chooses quality H when $p_1 < p \leq v$ and chooses quality L when $p_0 \leq p < p_1$ (and is indifferent to quality when $p = p_1$). The equilibrium expected profit of each firm is

$$\pi = \lambda(v-c)P^{N-1}\ , \tag{6}$$

[7] The left-hand side of (1) is increasing in P and ranges from 0 to $+\infty$ for $P \in [0, 1]$, while the right-hand side is positive and decreases in P. Therefore, there is a unique solution to (1) for $P \in [0, 1]$.

equilibrium total welfare is

$$W = (v - c)[\lambda(1 - P^N) + (1 - \lambda)(1 - P)^N] \ , \tag{7}$$

and equilibrium aggregate consumer surplus is

$$S = (v - c)[\lambda(1 - P^N) + (1 - \lambda)(1 - P)^N - \lambda N P^{N-1}] \ . \tag{8}$$

Proof. See Appendix B. ■

Theorem 1 describes a mixed strategy equilibrium in which firms sometimes offer a high-quality product for a high price and sometimes offer a low-quality product for a low price. For this mixed strategy equilibrium to exist it, is important that a firm cannot price discriminate between the two types of consumers, for instance, by offering a high-quality product aimed at the attentive consumers and a low-quality product aimed at the inattentive consumers. If firms could do this, they would compete separately in each market, and competition would force each type of product's price down to its cost, just like the asymmetric pure strategy equilibrium of Lemma 1. Thus, the model applies to situations in which a potential consumer approaches a firm (e.g., by phone) and obtains a quote. If instead, a firm publishes its offers (e.g., in its marketing material), inattentive consumers may be puzzled why some contracts have lower headline prices than others, and may investigate further why the contracts have different prices, and hence become more informed about quality. For this reason, a firm may prefer not to offer two contracts (so that inattentive consumers are not alerted).

The presence of inattentive consumers allows firms to earn positive profits, even though there is only one product which has positive social value. The fact that firms may cheat by cutting quality to exploit inattentive consumers allows a non-cheating firm to raise its price to attentive consumers. Expression (6) shows that in equilibrium a firm's profit is as if it extracts monopoly profit from each of the λ attentive consumers (equal to $v - c$ per consumer) whenever all its rivals choose to cheat (which occurs with probability P^{N-1}).

When there are at least four suppliers, there are multiple equilibria: the asymmetric pure strategy equilibrium in Lemma 1 and the symmetric mixed strategy equilibrium in Theorem 1. The former yields zero profit, whereas the latter (except in extreme cases) yields positive profit. Therefore, we might expect that firms will find a way to coordinate on the more profitable equilibrium. In addition, if there were some small cost of entering the market, then this would rule out the Bertrand-style equilibrium, and leave only the mixed strategy equilibrium.

To illustrate the mixed strategy equilibrium, consider the example with $N = 2, \lambda = \frac{1}{2}, c = 1$ and $v = 2$. Then Theorem 1 implies that

$$p_1 = \sqrt{2}\ ,\ P = p_0 = 2 - \sqrt{2}$$

and

$$F(p) = \begin{cases} 1 - \frac{p_0}{p} & \text{if } p_0 \leq p \leq p_1 \\ 1 - \frac{1}{2}P\frac{2-p}{p-1} & \text{if } p_1 \leq p \leq 2 \end{cases}$$

which is depicted in this figure.

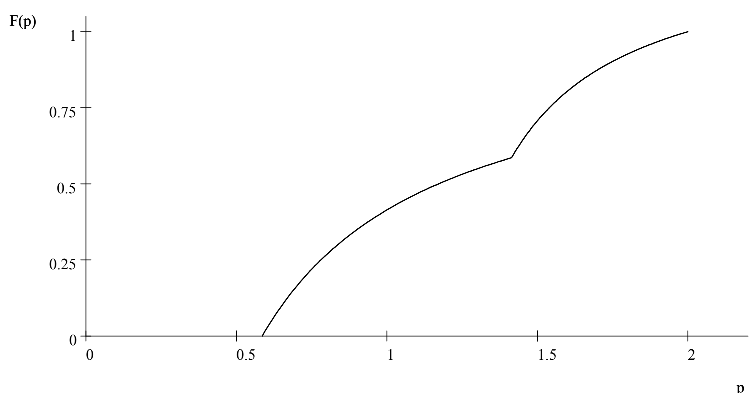

Figure 1: Cumulative distribution function for price p $(v = 2, c = 1, \lambda = \frac{1}{2}, N = 2)$

Attentive consumers buy the high-quality item whenever at least one firm supplies this item, which in this example occurs with probability $1 - P^2 \approx 0.66$. (If no firm supplies the high-quality item, attentive consumers do not buy at all.) Inattentive consumers buy the low-quality item whenever at least one firm supplies that item, which in this example occurs with probability $1 - (1 - P)^2 \approx 0.83$.

3 The Impact of Policy

In this section we investigate the impact on outcomes of changes in the market parameters λ and N. Changes in N may be the result of competition policy, which affects mergers and the ease of entry. Changes in λ may be induced by policy towards market transparency, for instance. Recall that in Varian's model, increases in λ benefit all consumers, while increases in N benefit attentive consumers but harm inattentive consumers. (Later in section 4.1 we discuss a third policy, which is to impose minimum quality standards on the market.)

3.1 Market transparency policy

Consider first the impact of λ on outcomes. When almost all consumers are attentive, so that $\lambda \approx 1$, one can check from (1) that the high-quality product is offered almost all the time ($P \approx 0$), and the equilibrium price is very likely to be close to cost c. The equilibrium profits are close to zero. When almost all consumers are inattentive ($\lambda \approx 0$), then firms cheat almost all the time ($P \approx 1$). Attentive consumers almost never buy, and firms compete to offer a low-quality product to only the inattentive consumers. The result is that the price is very likely to be close to zero. (In particular, $p_0 \approx 0$.)

The following result expands this discussion:

Proposition 1

(i) P decreases with λ, $P = 1$ when $\lambda = 0$ and $P = 0$ when $\lambda = 1$;

(ii) welfare W increases with λ, $W = 0$ when $\lambda = 0$ and $W = v - c$ when $\lambda = 1$;

(iii) profit π increases with λ when λ is small, decreases with λ when λ is large, and $\pi = 0$ if $\lambda = 0$ or $\lambda = 1$;

(iv) aggregate consumer surplus S decreases with λ when λ is small, increases with λ when λ is large, and $S = 0$ when $\lambda = 0$ and $S = v - c$ when $\lambda = 1$.

Proof. (i) Since the left-hand side of (1) is increasing in P and the right-hand side decreases with P and λ, it follows that P decreases with λ.

(ii) From (7), welfare is increasing in λ (for given P) and decreasing in P (for given λ). Hence

$$\frac{dW}{d\lambda} = \frac{\partial W}{\partial \lambda} + \frac{\partial W}{\partial P}\frac{dP}{d\lambda} > 0 \; .$$

(iii) For given λ, from (6) write $\pi(\lambda) = \lambda(v - c)(P(\lambda))^{N-1}$ to be equilibrium profit for each firm, where $P(\lambda)$ is the probability of cheating. Clearly, $\pi(\lambda) \geq 0$ for all λ. Also, since $P = 1$ when $\lambda = 0$ and $P = 0$ when $\lambda = 1$, we see that profit $\pi = \lambda\,(v - c)\,P^{N-1}$ is zero when $\lambda = 0$ or $\lambda = 1$. Therefore, π first increases in λ and eventually decreases in λ for $\lambda \in [0, 1]$.

(iv) From expression (8) we have $S = 0$ when $\lambda = 0$, and $S = v - c$ when $\lambda = 1$. Notice also that $S \leq v - c$ for all $0 \leq \lambda, P \leq 1$. In addition,

$$\begin{aligned} \frac{1}{v-c}\frac{dS}{d\lambda} &= (1-P^N) - (1-P)^N - NP^{N-1} \\ &\quad - \left[\lambda NP^{N-1} + (1-\lambda)\, N(1-P)^{N-1} + \lambda N\,(N-1)\,P^{N-2}\right]\frac{dP}{d\lambda}\ , \end{aligned}$$

which equals $-N$ when $\lambda = 0$. Therefore, S slopes down for sufficiently small λ.

Similarly, when $\lambda = 1$ S reaches its maximum value $v - c$, and hence is locally increasing at $\lambda = 1$. ■

These comparative statics results have both intuitive and surprising aspects. The cheating probability P is decreasing in the proportion of attentive consumers, as seems intuitive. This also implies that the average product quality in the market is higher when more consumers are attentive to quality, which implies that welfare increases when there are greater numbers of attentive consumers.

It is more intriguing that profit is non-monotonic in the consumer attentiveness. The interactions between firms' quality choices and prices appear to be the key to understanding the non-monotonic relationship between profit and attentiveness. The presence of inattentive consumers creates product differentiation for otherwise homogeneous producers, even though firms have the same equilibrium price/quality strategy. There is more product differentiation when λ is at some intermediate level.

It is also surprising that aggregate consumer surplus is non-monotonic in the attentiveness of consumers in the market. (Recall that in the Varian model, all consumers were better off when there was a greater proportion of attentive consumers present in the market.) When λ is small, increasing the proportion of attentive consumers actually lowers aggregate consumer surplus. When λ is small, an increase in λ reduces cheating (increases quality) in the market, but it also increases prices, and the latter effect dominates when λ is small. Thus, when most consumers are inattentive, a market transparency policy which increases λ will harm the inattentive consumers (but always help the attentive consumers). When λ is large, further increases in λ not only increase quality in the market, but may also lower prices. These two effects are likely to work in the same direction for large λ, so that consumer surplus increases in λ when λ is large. In an example aggregate consumer surplus as a function of λ looks like the thick line on

Figure 2:

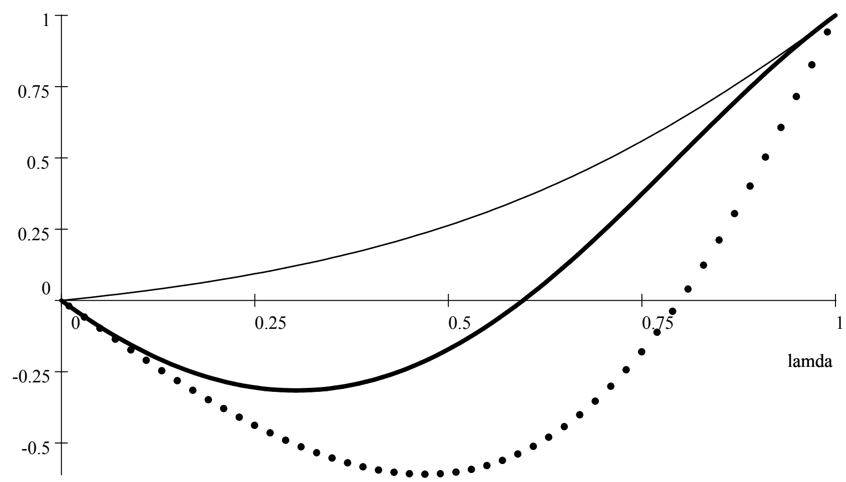

Figure 2: Consumer surplus as a function of λ $(N = 2, v = 2, c = 1)$

Of course, the impact differs on the two consumer groups. For instance, the surplus obtained by attentive consumers can never be negative. The thin solid line shows the consumer surplus (per consumer) enjoyed by the attentive consumers, while the dotted line shows the per-consumer surplus of the inattentive consumers. In particular, when there are relatively few inattentive consumers, those consumers obtain a positive surplus since firms so rarely cheat. Notice that attentive consumers are harmed by the presence of inattentive consumers. (If there were no inattentive consumers then attentive consumers would obtain the maximum possible surplus, $v - c$.) This contrasts with the models of "add-on" pricing, where the sophisticated consumers can be cross-subsidised by the high prices the naive consumers end up paying (see Ellison (2005) and Gabaix & Laibson (2006)).

3.2 Competition policy

Consider next the impact of N on market outcomes. This turns out to be less clear-cut than the impact of λ, and we do not describe the comparative statics in all situations. Expression (1) for the cheating probability could easily be used in the previous section to obtain results on the effect of λ, but it is less useful to generate corresponding results for changes in N. For that purpose, the analysis can be simplified by using an approximation for the cheating probability. Whenever the term $\frac{1-\lambda}{\lambda}(1-P)^{N-1}$ is small, expression (1) indicates that

$$P \approx \frac{k^{\frac{1}{N-1}}}{k^{\frac{1}{N-1}}+1}\ , \text{ where } k = \frac{1-\lambda}{\lambda}\frac{c}{v-c}\ . \tag{9}$$

This approximation is accurate whenever λ is close to 1 or N is large. In this section we focus on cases where N or λ is reasonably large so that we can use (9).

First, it is straightforward to characterize the limit outcome for large N. From (1) we have

$$\frac{P}{1-P} = \left\{k\left[1+\frac{1-\lambda}{\lambda}(1-P)^{N-1}\right]\right\}^{\frac{1}{N-1}}\ ,$$

where k is given in (9). Since the term in brackets $\{\cdot\}$ is positive and bounded above by k/λ and bounded below by k, the limit of the right-hand side of the above as $N \to \infty$ always exists and is equal to 1. It follows that

$$P \to \frac{1}{2} \text{ as } N \to \infty\ .$$

(This can also be seen in the approximation (9).) Thus we see that, regardless of parameter values for v, c and λ, in markets with many firms, each firm chooses to cheat approximately half the time!

Since $P \approx \frac{1}{2}$ for large N, expression (6) implies that both an individual firm's profit and industry profit converge to zero as the number of firms becomes large. (This contrasts with Varian's model, where industry profit did not depend on the number of firms.) Similarly, (7) implies that welfare converges to $\lambda(v-c)$ for large N. This in turn implies that aggregate consumer surplus also converges to $\lambda(v-c)$. The impact on the two types of consumer is a little more involved, but still clear-cut. From (2) we see that $p_1 \to c$ and from (3) we see that $p_0 \to 0$, both as $N \to \infty$. Then (4) and (5) imply that the distribution for prices converges to a discrete two-point distribution, with probability half on $p = 0$ and probability half on $p = c$. Therefore, the surplus of each of the attentive consumers converges to $v - c$, while the surplus of each of the inattentive consumers converges to zero. This limit outcome corresponds exactly to the asymmetric pure strategy equilibrium described in Lemma 1.

We next consider how changes in N affects the cheating probability, industry profit, welfare, and aggregate consumer surplus:

Cheating probability: If N or λ is relatively large, we can use the formula (9) for the cheating probability P. Clearly, P in (9) increases (decreases) with N if $k < 1$

(respectively, $k > 1$). Note that $k < 1$ is equivalent to the condition

$$\lambda > \frac{c}{v} \ .$$

In particular, if the fraction of attentive consumers is large enough, the cheating probability starts small and increases (with asymptotic limit $\frac{1}{2}$) as the number of sellers increases. If we interpret P as the degree of ethical behavior in the market, then increased competition, in the sense of having more firms in the market, leads to less ethical behavior whenever there is a small fraction of inattentive consumers.[8]

Industry Profits: It turns out that industry profits can rise or fall with the addition of a further seller. For instance, for fixed $\lambda \in (0,1)$, when N is large the approximation (9) is valid, and industry profit is proportional to

$$\frac{N}{(1+k^{\frac{1}{N-1}})^{N-1}} \approx \frac{N}{2^{N-1}} \ .$$

This is decreasing in N for large N. By contrast, for fixed N (not necessarily large), when $\lambda \approx 1$ the approximation (9) is valid, and the impact on industry profit of introducing one more seller is proportional to

$$\frac{N+1}{(1+k^{\frac{1}{N}})^{N}} - \frac{N}{(1+k^{\frac{1}{N-1}})^{N-1}} \ .$$

This term is positive when $\lambda \approx 1$ (i.e., when $k \approx 0$). Thus, in the "almost rational" case where only a small fraction of consumers are inattentive, the addition of more sellers boosts industry profit.

Welfare: In a similar way, it is possible to characterize the impact of an additional seller on welfare when N is large or λ is close to 1. Using the approximation (9), we have

$$\frac{W}{v-c} \approx \lambda + \frac{1-\lambda-\lambda k^{\frac{N}{N-1}}}{\left(1+k^{\frac{1}{N-1}}\right)^{N}} \ .$$

Therefore, the differential of W with respect to N has the same sign as

$$\frac{dW}{dN} \overset{\text{sign}}{=} (\ln k)\, \lambda \frac{k^{\frac{N}{N-1}}}{(N-1)^2} +$$

[8]See Shleifer (2004) for further discussion about the effect of competition on ethical behavior.

$$\left(1-\lambda-\lambda k^{\frac{N}{N-1}}\right)\left(N\left(\ln k\right)\frac{k^{\frac{1}{N-1}}}{(N-1)^2\left(k^{\frac{1}{N-1}}+1\right)}-\ln\left(k^{\frac{1}{N-1}}+1\right)\right).$$

In particular, for given $\lambda \in (0,1)$, as N becomes large the only term in the above which does not vanish is

$$-\left(1-\lambda-\lambda k^{\frac{N}{N-1}}\right)\ln\left(k^{\frac{1}{N-1}}+1\right)\approx(\lambda k-(1-\lambda))\ln 2\ .$$

Therefore, for large N welfare increases (decreases) with N when $\lambda k > 1-\lambda$ (respectively, $\lambda k < 1-\lambda$). Note that $\lambda k > 1-\lambda$ if and only if $2c > v$ (which does not depend on λ).

Similarly, for fixed N, as $\lambda \to 1$ we can see that dW/dN is negative.[9] Thus, in markets where only a few consumers are inattentive, adding more competitors is bad for welfare. Intuitively, when λ is close to 1, all that matters for welfare in (7) is the payoff to the attentive consumers, which is proportional to $1-P^N$. One can show that P^N is increasing in N whenever λ is sufficiently close to 1, and hence attentive consumers are harmed by the addition of a further seller.

Consumer surplus: Again, if either N is large or λ is close to 1, from (8) we have

$$\frac{S}{v-c}\approx\lambda+\frac{1}{\left(1+k^{\frac{1}{N-1}}\right)^N}\left[1-\lambda-\lambda k^{\frac{N}{N-1}}\right]-\lambda\frac{Nk}{\left(1+k^{\frac{1}{N-1}}\right)^{N-1}}$$

$$=\lambda+\frac{1-\lambda-\lambda k^{\frac{N}{N-1}}-\lambda Nk(1+k^{\frac{1}{N-1}})}{\left(1+k^{\frac{1}{N-1}}\right)^N}\ .$$

Therefore, the differential of S with respect to N has the sign of

$$\frac{dS}{dN}\overset{\text{sign}}{=}\left(k\lambda\left(-k^{\frac{1}{N-1}}-1\right)+(\ln k)\,\lambda\frac{k^{\frac{N}{N-1}}}{(N-1)^2}+N\left(\ln k\right)\frac{\lambda}{(N-1)^2}k^{\frac{N}{N-1}}\right)+$$

$$\left(Nk\lambda\left(-k^{\frac{1}{N-1}}-1\right)-\lambda k^{\frac{N}{N-1}}-\lambda+1\right)\left(N\left(\ln k\right)\frac{k^{\frac{1}{N-1}}}{(N-1)^2\left(k^{\frac{1}{N-1}}+1\right)}-\ln\left(k^{\frac{1}{N-1}}+1\right)\right).$$

For fixed $\lambda \in (0,1)$ one can show this expression is positive for sufficiently large N. (As N becomes large in the above, there is just one explosive term, which is positive, and this is $-Nk\lambda\left(-k^{\frac{1}{N-1}}-1\right)\ln\left(k^{\frac{1}{N-1}}+1\right)$.) In effect, welfare converges to its limit much

[9] The easiest way to see this is to divide the expression for dW/dN by $k^{\frac{N}{N-1}}$ and let k tend to zero.

faster than profit converges to zero, and consumers and firms end up playing a constant-sum game when N is reasonably large. Since firms' profits fall with N, it follows that consumers must gain. Finally, for fixed N (not necessarily large), consumers in aggregate are worse off with the introduction of an additional seller whenever λ is close enough to 1. Since we know that welfare decreases with N in this case, while industry profit rises, it necessarily follows that consumer surplus (equal to welfare minus profit) must fall.

We summarize this discussion in the following two propositions:

Proposition 2 *For fixed $\lambda \in (0,1)$ and N sufficiently large:*
(i) the outcome converges to the asymmetric pure strategy equilibrium described in Lemma 1, namely: industry profit is zero; welfare and aggregate consumer surplus is $\lambda(v-c)$, and inattentive consumers obtain zero surplus;
(ii) each firm cheats with probability approximately equal to $\frac{1}{2}$;
(iii) the cheating probability decreases with N if $\lambda < c/v$ and increases with N if $\lambda > c/v$;
(iv) industry profit falls with N;
(v) aggregate consumer surplus rises with N; and
(vi) welfare falls with N if $v > 2c$ and rises with N if $v < 2c$.

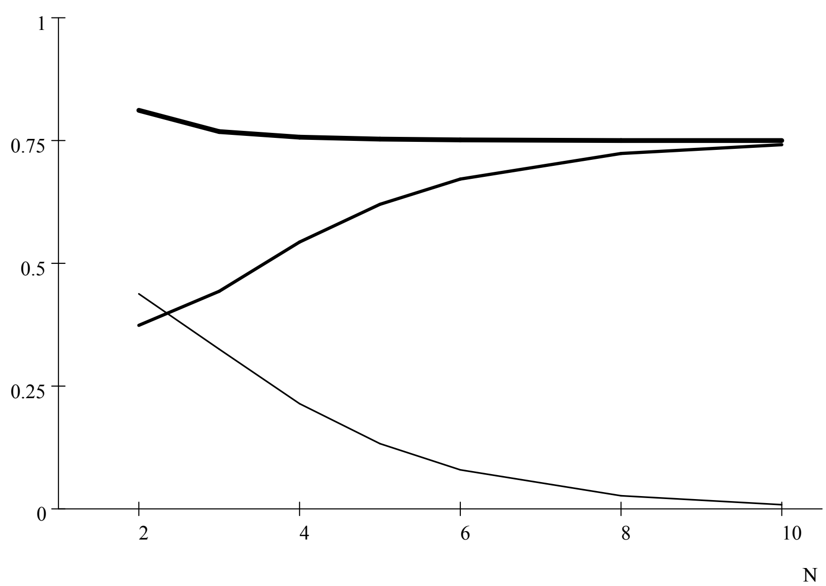

Figure 3: Welfare (thick line), industry profit (thin line) and aggregate consumer surplus (medium line) as the number of firms changes ($\lambda = 0.75$, $v = 2, c = 1$)

These results are illustrated in Figure 3 (which uses the exact formula for the cheating probability in (1) rather than the approximation (9). Here, welfare decreases slightly with N but converges quickly to its limit, while profits fall and consumer surplus rises monotonically with N. Notice that N does not need to be especially large for the results reported in Proposition 2 to be observed.

Next, we summarize the results when most consumers are attentive:

Proposition 3 *Fix the initial number of firms N. If λ is sufficiently close to 1 then when an additional seller is introduced:*
(i) the cheating probability rises;
(ii) industry profit rises;
(iii) welfare falls;
(iv) aggregate consumer surplus falls.

We should not place too much weight on the results in Proposition 3, which are somewhat pathological. Numerical simulations suggest that the fraction of inattentive consumers needs to be *extremely* low in order for, say, profit to rise when there are more sellers. In any case, when $\lambda \approx 1$, profit, welfare and consumer surplus are all essentially flat with respect to the number of sellers, and the impact (positive or negative) of changing the number of sellers is negligible.

We have not yet mentioned the impact of more competition on the two kinds of consumer separately. In general, this is a difficult task. However, it is straightforward to compare the outcomes in the two extreme cases of $N = 2$ and $N = \infty$. With duopoly, the outcomes for the two kinds of consumer are illustrated on Figure 2 above (in the particular case with $v = 2$ and $c = 1$). With many firms, the outcome for attentive consumers is that they each have surplus $v - c$ while the outcome for inattentive consumers is that they have zero surplus. Comparing the two extreme market structures shows that the attentive consumers are always better off in competitive markets compared to duopoly (as in the Varian model). The impact on inattentive consumers depends upon how numerous they are. When λ is relatively large, these consumers can "free ride" on the attentive consumers in the duopoly setting and obtain positive surplus, which is eroded when there are more suppliers. Thus, in this case increased competition harms the inattentive consumers (as in the Varian model). But if λ is relatively small, these consumers are strictly exploited with duopoly, whereas they are simply given nothing when there are many sellers. More competition helps these consumers in this case (unlike in Varian's model).

4 Discussion and Extensions

4.1 Heterogeneous taste for quality

Varian's model, in which some consumers do not search beyond their initial random choice of supplier, has two interpretations. A "rational" interpretation of the uninformed consumers is that they incur extremely high search costs to sample a second seller. A "boundedly rational" interpretation is that these consumers mis-perceive how well the market operates, and mistakenly believe it is perfectly competitive in the sense that all sellers offer the same price; in this situation there is perceived to be no point in sampling another seller.

Like Varian's model, our model also has a "rational" interpretation: all consumers have perfect information about prices and qualities available in the market, but a fraction $1-\lambda$ of them do not suffer utility loss from consuming the low-quality product (their utility is not sensitive to quality). For instance, some consumers value good quality wine, while others will just buy the cheapest bottle they can find. Our model would then be one of vertical product differentiation with price and quality dispersion. Firms' strategies, described in Theorem 1, would be unchanged in this alternative interpretation. Our results about how changes in λ and N affected product quality, firm profits and the surplus of the attentive (or quality-sensitive) consumers would continue to be valid too. However, this interpretation would change our results concerning the surplus of the inattentive (or quality-insensitive) consumers, and it would also change the kinds of markets to which the model might apply.

Returning to the interpretation in which some consumers are inattentive to quality (rather than simply not valuing quality), our main model is restrictive in its assumption that all consumers have the same tastes for quality. In particular, our model exhibits market failure which can easily be corrected by means of a simple policy, which is to impose minimum quality standards and force all firms to offer the high-quality product. Since all consumers value the high-quality product but do not value the low-quality product, a policy which forbids supply of the low-quality product—assuming this can be enforced—will lead to the ideal outcome: all consumers buy the high-quality item and its price is equal to cost. However, this simple "one-size-fits-all" policy remedy is an artifact of our assumption that all consumers have the same preferences for quality. A richer model would allow some consumers to prefer the low-quality product (with a lower price) to the high-quality product. In this case there are two reasons why a low-quality product is chosen: a consumer does not pay attention to quality, or the consumer is careful but prefers this product due to a low willingness-to-pay for quality. In this richer framework where some attentive consumers actually prefer the low-quality product at its lower price, the impact of imposing minimum quality standards is less clear cut: a ban on the low quality product will help the inattentive consumers but harm those consumers who actively want the low-quality item.[10]

To discuss this important trade-off in more detail, consider this variant of our main model.[11] The product can be supplied with two levels of quality, q_L and q_H, and the respective unit costs of supplying this product are c_L and c_H. For reasons of tractability, suppose there is a perfectly competitive market ($N = \infty$), and each variety is available for a price equal to its cost. (Alternatively, there may be a finite number of firms, but the asymmetric equilibrium described in Lemma 1 is chosen.) In terms of preferences, there are two kinds of consumers (unlike our main model): those who value the high-quality item highly, and those who do not. Specifically, the consumers who value high quality have utility $\theta_H q - p$ if they consume a product with quality q and price p, while the other consumers have utility $\theta_L q - p$ (where $\theta_H > \theta_L$) for the same product. Suppose it is efficient for the type-θ_H consumers to buy the high-quality item and for the type-θ_L consumers to buy the low-quality item, i.e.,

$$\theta_H q_H - c_H > \theta_H q_L - c_L \; ; \; \theta_L q_L - c_L > \theta_L q_H - c_H \; . \tag{10}$$

Suppose that a fraction α of consumers have taste parameter θ_H. (Our main model essentially assumed $\alpha = 1$.)

As well as possessing these taste differences, consumers differ in their attention to quality. As before, a fraction $1-\lambda$ of type-θ_H consumers do not think about quality when they decide on their product, and buy simply on the basis of price. (It does not matter

[10] A similar trade-off is seen in Leland (1979). Leland analyzes a lemons-type situation where (all) consumers can observe the average quality in a market, but not the quality of an individual item. A policy intervention which raises the minimum quality of products in the market will help overcome the lemons problem, but will prevent some consumers buying the (low-quality) items they want.

[11] This discussion is taken from Armstrong (2008).

if the type-θ_L consumers think about quality or not, as they will buy the appropriate product even if they buy only on the basis of price.) Without policy intervention, all type-θ_L consumers buy the low-quality product for a price c_L, as do that fraction of type-θ_H consumers who do not pay attention to quality. The remaining type-θ_H consumers buy the high-quality product for price c_H. Therefore, welfare without intervention is

$$W = (1-\alpha)\left[\theta_L q_L - c_L\right] + \alpha\left[\lambda(\theta_H q_H - c_H) + (1-\lambda)(\theta_H q_L - c_L)\right] \ .$$

On the other hand, suppose policy prevents the supply of the low-quality product in order to protect those consumers who value quality but do not pay attention to it. In this case, on the assumption that the type-θ_L consumers prefer buying the high-quality product to buying nothing, welfare when choice is restricted is

$$W = (1-\alpha)[\theta_L q_H - c_H] + \alpha[\theta_H q_H - c_H] \ .$$

Welfare is increased by imposing a minimum quality standard whenever

$$(1-\alpha)[\theta_L(q_H - q_L) - (c_H - c_L)] + \alpha(1-\lambda)[\theta_H(q_H - q_L) - (c_H - c_L)] > 0 \ .$$

From (10), the first term in square brackets is negative while the second term in square brackets is positive. Thus, whether this policy intervention improves welfare depends on the relative sizes of α and λ, as well as the magnitude of the loss incurred by those consumers who actively want the low-quality item and the magnitude of the gain experienced by those inattentive consumers who like high quality. Except in extreme cases, the informational requirements needed to be confident that the interventionist, choice-restricting policy is desirable are substantial.

4.2 When beliefs are Bayesian

An important assumption in our main model (with homogeneous tastes for quality) is that the uninformed consumers do not attempt to infer a seller's likely product quality from the seller's price. When consumers buy the product repeatedly, this assumption may be implausible. What happens if the uninformed consumers have rational expectations about the equilibrium relationship between price and quality in the market? The papers closest to ours in modelling terms, but where the consumers who are not directly informed about quality are Bayesian, are probably Wolinsky (1983) and Cooper & Ross

(1984). Wolinsky's model allows consumers to have heterogeneous tastes for quality, and to obtain noisy information about a seller's quality whenever a buyer considers buying from a seller. He finds conditions under which a separating equilibrium exists, so that consumers can infer quality from a firm's price in equilibrium. Prices must be above costs in order that a seller does not have an incentive to cheat on quality. If consumer information about quality is good, the price-cost markup need not be large, but in other situations prices need to be substantially above cost to prevent a firm from cutting quality.

Cooper & Ross present a model (like ours) in which consumers have homogenous preferences over quality, and are exogenously divided into a group which accurately knows product quality and a group which has no direct knowledge of quality. Unlike our model, though, the latter group has rational expectations about the relationship between price and quality in the market. In addition, unlike our model, Cooper & Ross focus on perfectly competitive markets. They find that no rational expectations equilibrium exists when firms have constant returns to scale (as they do in our model), but when firms have U-shaped average costs equilibria can exist in which all firms charge the same price (so no inference about quality from price is possible) but where both high and low quality products are sold.

To illustrate some of the ideas from these two earlier papers, consider how Bayesian consumers can sometimes support a pooling equilibrium in which only the high-quality product is supplied. (Except for the fact that the uniformed consumers are Bayesians, the framework is unchanged from our previous model.) Suppose there are two suppliers, and let p^* be a candidate price at which both firms offer the high-quality product for sure. Suppose that those consumers who do not directly observe quality assume that any firm which offers a lower price than p^* is offering a low-quality product. For this situation to be an equilibrium, two conditions must be met. First, a firm which offers the price p^* but cheats and offers a low-quality product should not make greater profit, i.e., we require

$$\tfrac{1}{2}(p^* - c) \geq \tfrac{1}{2}(1 - \lambda)p^* \ . \tag{11}$$

(The left-hand side is a firm's profit if it follows the equilibrium strategy, while the right-hand side is the firm's profit if it instead offers a low-quality product, in which case it loses all its attentive consumers but has a lower cost.) Second, the firm which continues to offer a high-quality product but slightly undercuts its rival in terms of price cannot make more profit, i.e.,

$$\tfrac{1}{2}(p^* - c) \geq \lambda(p^* - c) \ . \tag{12}$$

(In the deviation, the firm loses its inattentive consumers who presume the firm now offers a low-quality product, but it attracts all the attentive consumers.) Condition (12) requires $\lambda \leq \frac{1}{2}$, so that only a minority of consumers can be attentive. Condition (11) requires that the pooling price be sufficiently high. Since the maximum feasible price is v, a price p^* can be found which satisfies (11) whenever $\lambda \geq c/v$.

In summary, whenever parameters are such that

$$\frac{c}{v} \leq \lambda \leq \tfrac{1}{2} \ ,$$

a continuum of pooling equilibria exist with price p^* in the interval

$$\frac{c}{\lambda} \leq p^* \leq v$$

in which both firms offer the high-quality product for sure. Thus, unlike with our main model, there is no market failure even though a fraction of consumers do not directly observe quality. In this equilibrium, price exceeds cost in order to prevent firms being tempted to cheat on quality. However, when there are more suppliers in the market, these pooling equilibria exist only for a narrower range of parameter values. When a firm has a small market share, the temptation to undercut the market price to attract all attentive consumers becomes stronger. Indeed, in the competitive limit as the number of firms becomes large, no pooling equilibrium exists. (This is essentially Proposition 2 in Cooper & Ross (1984).)

TECHNICAL APPENDICES

APPENDIX A: Varian's Model of Inattention to Price

Recall Varian's (1980) model of sales. There, a population of consumers each wish to buy a single unit of a homogeneous product. There are N identical risk-neutral firms in the market, each with marginal cost c. A fraction λ of consumers know all the prices in the market and buy from the lowest-price supplier. The remaining $1 - \lambda$ consumers buy from a random supplier, so long as that supplier's price is no greater than the reservation utility $v > c$. Each consumer has a unit demand, and the consumer population size is normalized to 1.

In this environment there can be no pure strategy equilibrium. For instance, if a firm knows that each of its rivals is charging price $p > c$ for sure, it can slightly undercut this price and so attract the entire pool of attentive consumers. Or if the other firms set a price equal to cost, a firm can make positive profit by exploiting its captive, inattentive consumers, who are $(1 - \lambda)/N$ in number. Therefore, firms follow a mixed pricing strategy.

A symmetric mixed pricing strategy will involve each firm's price being taken from the cdf $F(p)$, where F satisfies

$$\left[\lambda\left(1 - F(p)\right)^{N-1} + \frac{1-\lambda}{N}\right](p - c) \equiv \frac{1-\lambda}{N}(v - c) \ . \tag{13}$$

The term in square brackets represents a firm's expected demand when it sets price p: it will attract the λ attentive consumers whenever it sets the lowest price, which occurs with probability $(1 - F(p))^{N-1}$; and the firm will always attract its share of the inattentive consumers, $(1 - \lambda)/N$, as long as $p \leq v$. Therefore, the left-hand side is a firm's expected profit with the price p. The firm must be indifferent between all prices in the support of F. One can show that $p = v$ is in the support, in which case a firm's profit is always equal to the right-hand side. Expression (13) can be rearranged to give the explicit solution

$$F(p) = 1 - \left(\frac{1-\lambda}{N\lambda}\frac{v-p}{p-c}\right)^{\frac{1}{N-1}} \ .$$

The lower end of the support, p_0, is then the price which makes $F(p) = 0$.

One can calculate the expected prices paid by the attentive and the inattentive consumers, and see how these prices depend on the main parameters of interest, N and λ. As the number of suppliers becomes large, one can show that the attentive consumers pay an expected price close to cost, while the inattentive consumers pay an expected price close to the monopoly price v. This is depicted in Figure 4 for a particular example. Competition policy—which acts to forbid mergers or ease the entry of new firms—will therefore help attentive consumers (whose expected price is the thin solid line on the figure) but *harms* the inattentive consumers (whose expected price is the dotted line). Notice that industry revenue (indicated by the thick line on the figure) does not depend on the number of suppliers, since each firm's profit is equal to its share of the captive market. Industry profit is always $(1 - \lambda)(v - c)$ in this model.

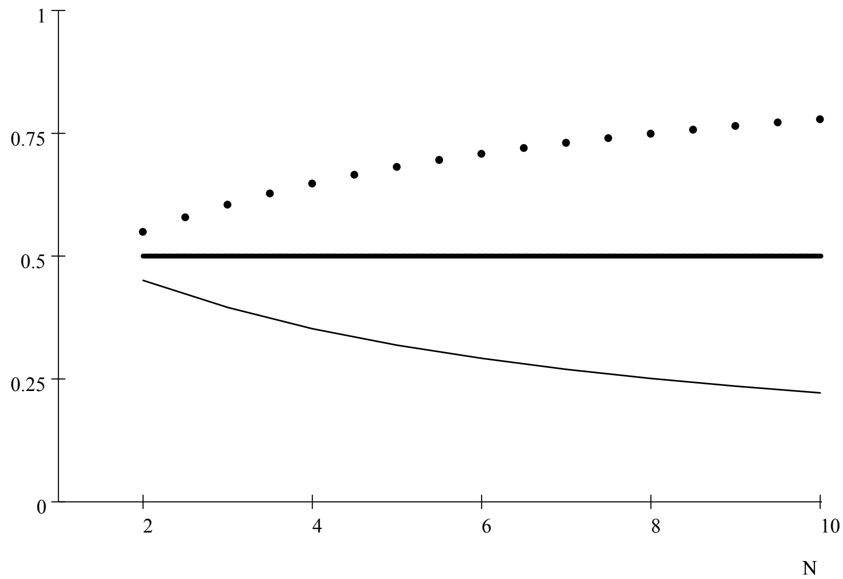

Figure 4: Expected prices paid by attentive and inattentive consumers for $2 \leq N \leq 10$ $(v = 1,\ c = 0,\ \lambda = \frac{1}{2})$

Turning next to the impact of λ on market prices, consider Figure 5.

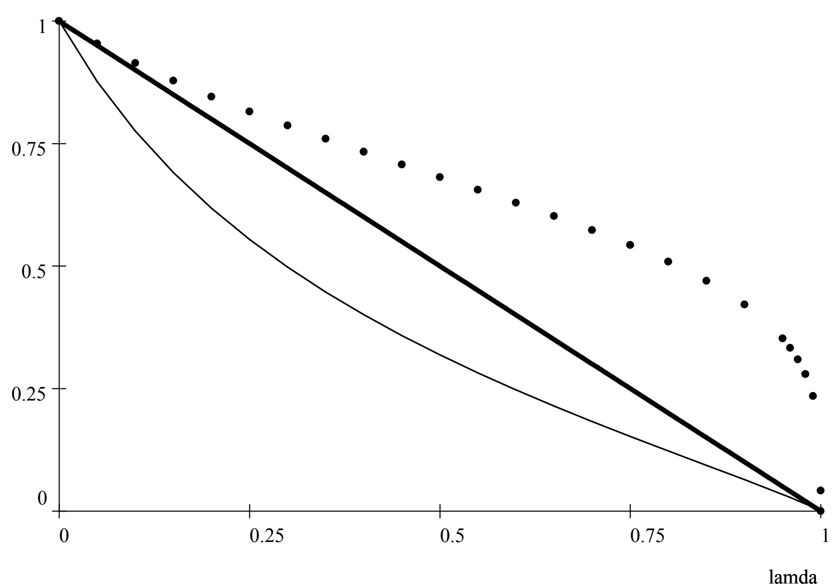

Figure 5: Expected prices paid by attentive and inattentive consumers as a function of λ $(N = 5,\ v = 1,\ c = 0)$

Here, the thick line represents industry revenue (which is linear in λ), the thin solid line depicts the price paid by attentive consumers, and the dotted line represents the price paid by inattentive consumers. Increasing the fraction of attentive consumers benefits both groups of consumers, and reduces industry profit. Total welfare—the sum of aggregate consumer surplus and industry profit—is equal to $v - c$ in Varian's model and depends neither on N nor λ.

APPENDIX B: Proof of Theorem 1

First, we establish that $F(\cdot)$ described in the statement of Theorem 1 is a continuous cdf with support $[p_0, v]$. Clearly, F in (4) is increasing in p and satisfies $F(p_0) = 0$. Similarly, F defined (implicitly) by (5) is an increasing function of p, where $F(p)$ ranges from P to 1 as p ranges from p_1 to v. (As $F(p)$ on the left-hand side of (5) ranges from P to 1, the p which makes the right-hand side equal ranges from p_1 (where this follows from the definition of p_1 in (2)) to v.) Finally, we need to show that $F(\cdot)$ is continuous at the point p_1 which separates expressions (4) and (5). As we have said, if we set $p = p_1$ in expression (5), the definition of p_1 in (2) implies that $F(p_1) = P$. And if we set $p = p_1$ into expression (4), from the definition of p_0 in (3) we obtain that $F(p_1) = P$. Thus, the two parts of $F(\cdot)$ join up appropriately, and F is indeed a cdf.

Suppose $N - 1$ firms follow the strategy described. If the remaining firm chooses a price p with $p_0 \leq p \leq p_1$ and quality L, its expected profit is

$$(1 - \lambda)\, p \left[1 - F(p)\right]^{N-1} .$$

(The firm will sell only to inattentive consumers, and then only if it sets the lowest price, which occurs with probability $(1 - F(p))^{N-1}$.) A firm is indifferent between any $p \in [p_0, p_1]$ if

$$(1 - \lambda)\, p \left[1 - F(p)\right]^{N-1} = (1 - \lambda)\, p_0 = (1 - \lambda)\, p_1 (1 - P)^{N-1} ,$$

where if a firm sets p_0, it will surely sell to inattentive consumers, and if it sets p_1 (and offers quality L), it will sell to inattentive consumers whenever all other firms set a higher price. The first equality above implies (4), while the second equality implies (3). Since $0 \leq P \leq 1$ in (1), it follows that $p_0 \leq p_1$ as required.

Alternatively, if a firm chooses a price p with $p_1 \leq p \leq v$ and quality H, its expected profit is

$$(p - c) \left\{ (1 - \lambda) \left[1 - F(p)\right]^{N-1} + \lambda \left[P + 1 - F(p)\right]^{N-1} \right\} , \tag{14}$$

where if a firm charges $p \geq p_1$ and sets quality H, it sells to inattentive consumers only if it has the lowest price, and it sells to the attentive consumers if no other firm sets a price between p_1 and p. The probability that no other firm sets a price between p_1 and $p \geq p_1$ is

$$[F(p_1) + 1 - F(p)]^{N-1} = [P + 1 - F(p)]^{N-1} \quad .$$

A firm is indifferent between any $p \in [p_1, v]$ if expression (5) holds, where if a firm charges v it guarantees itself an expected profit $\lambda (v - c) P^{N-1}$ (i.e., it sells to the attentive consumers at the monopoly price when the other firms have all cheated and chosen quality L).

So far, all of this analysis is valid for arbitrary $0 \leq P \leq 1$. The equilibrium choice of P is determined by the requirement that a firm be indifferent between being a low-price/low-quality firm and being a high-price/high-quality firm. This is ensured if at the point $p = p_1$ a firm is indifferent between offering a low or a high quality product. At this price, if a firm decides to offer a high-quality product its demand jumps (it then attracts all the attentive consumers), but its unit cost jumps too, from zero to c. For a firm to be indifferent, these two jumps must cancel out, so that

$$\underbrace{p_1(1-\lambda)(1-P)^{N-1}}_{\text{profit with low quality}} = \underbrace{(p_1 - c)(\lambda + (1-\lambda)(1-P)^{N-1})}_{\text{profit with high quality}} . \tag{15}$$

However, expression (15) when combined with (2) yields the formula (1) for P. In sum, a firm is indifferent between all strategies of the forms: (i) set price $p_0 \leq p \leq p_1$ and offer low quality and (ii) set price $p_1 \leq p \leq v$ and offer high quality. In all cases, the firm makes the profit obtained by choosing the highest price $p = v$ and offering the high-quality product, which yields profit (6).

To complete our proof that the candidate equilibrium is indeed an equilibrium, we need to show that the remaining deviations are not profitable, and these deviations are (i) setting price $p_0 \leq p < p_1$ and high quality and (ii) setting price $p_1 < p \leq v$ and low quality. (If these deviations are not profitable a firm also cannot benefit from a randomized deviation in quality choice.) If a firm chooses price $p_0 \leq p < p_1$ and high quality, its profit is

$$(p - c)\left\{(1-\lambda)\left[1 - F(p)\right]^{N-1} + \lambda\right\} = (p - c)\left\{(1-\lambda)\frac{p_0}{p} + \lambda\right\},$$

which is increasing in p within this range. (The above equality follows from (4).) Therefore, the highest profit the firm can obtain using this form of strategy involves choosing

$p = p_1$, in which case the firm obtains just the equilibrium profit π in (6). Therefore, this deviation cannot be profitable.

The argument for deviation (ii) is more complicated. If a firm sets price $p_1 < p \leq v$ and low quality, its profit is

$$(1-\lambda)\, p\, [1 - F(p)]^{N-1} \ .$$

From (5), if the cdf at the price $p \geq p_1$ is $F \geq P$, the price in terms of F is

$$p = c + \frac{\lambda(v-c)P^{N-1}}{(1-\lambda)(1-F)^{N-1} + \lambda(P+1-F)^{N-1}} \geq p_1 \ .$$

Therefore, the firm's profit in terms of F instead of p is

$$(1-\lambda)(1-F)^{N-1}\left[c + \frac{\lambda(v-c)P^{N-1}}{(1-\lambda)(1-F)^{N-1} + \lambda(P+1-F)^{N-1}}\right] \ . \qquad (16)$$

By construction this profit equals π in (6) when $F = P$. However, one can show that (16) is decreasing in F, and so the firm cannot strictly gain by such a deviation.

Welfare is zero whenever a low quality product (or no product) is supplied to a consumer; otherwise welfare is $v - c$ for each consumer served. All λ attentive consumers are served with high quality when at least one firm does not cheat, which occurs with probability $(1 - P^N)$. The $1 - \lambda$ inattentive consumers are served with high quality only when no firm cheats, which occurs with probability $(1 - P)^N$. This explains expression (7). Since aggregate consumer surplus is the difference between welfare and industry profit, (8) follows from (6) and (7).

Finally, we show that there can be no other symmetric equilibrium. From Lemma 1, there can be no pure strategy symmetric equilibrium, so suppose that there is another mixed-strategy symmetric equilibrium. We only need to consider strategies in which prices follow some cdf $G(p)$ and quality choices are some randomization between H and L for any p. Denote the lower and upper support of G by p_l and p_h, respectively. Then $0 < p_l < p_h \leq v$, and each firm's expected profit is positive. At such an equilibrium, it is not possible that each firm always chooses quality H or each firm always chooses quality L. Notice that at p_h, each firm must choose H, otherwise a firm's profit at p_h would be zero. Also, $G(p)$ cannot have a mass point at p_h. Thus there exists some $p'' < p_h$ such that each firm chooses H for $p \geq p''$.

If for $p \in [p_l, p'']$ each firm chooses L, then the equilibrium must be the same as the one we have already characterized. Otherwise, there exists some $p' < p''$ such that each firm chooses L for $p \in (p', p'')$ and chooses H at $p' > c$, where $G(p)$ is continuous at p' and p''. Then, at prices p' and p'', the expected profit from choosing L must be the same, or

$$p'(1-\lambda)[1-G(p')]^{N-1} = p''(1-\lambda)[1-G(p'')]^{N-1}.$$

Similarly, at prices p' and p'', the expected profit from choosing H must also be the same, or

$$(p'-c)\left[(1-\lambda)[1-G(p')]^{N-1}+\lambda\gamma\right] = (p''-c)\left[(1-\lambda)[1-G(p'')]^{N-1}+\lambda\gamma\right],$$

where γ is the probability that each firm will sell to attentive consumers at p''. Notice that since all firms choose L for $p' < p < p''$, the probability that each firm will sell to attentive consumers at p' is equal to that at p''. It follows that

$$\frac{p''}{p'} = \frac{[1-G(p')]^{N-1}}{[1-G(p'')]^{N-1}}$$

and

$$\frac{p''-c}{p'-c} = \frac{(1-\lambda)[1-G(p')]^{N-1}+\lambda\gamma}{(1-\lambda)[1-G(p'')]^{N-1}+\lambda\gamma}.$$

Since $[1-G(p')] > [1-G(p'')]$ and $\lambda\gamma > 0$, we have

$$\frac{(1-\lambda)[1-G(p')]^{N-1}+\lambda\gamma}{(1-\lambda)[1-G(p'')]^{N-1}+\lambda\gamma} < \frac{(1-\lambda)[1-G(p')]^{N-1}}{(1-\lambda)[1-G(p'')]^{N-1}} = \frac{[1-G(p')]^{N-1}}{[1-G(p'')]^{N-1}} = \frac{p''}{p'}.$$

On the other hand,

$$\frac{(1-\lambda)[1-G(p')]^{N-1}+\lambda\gamma}{(1-\lambda)[1-G(p'')]^{N-1}+\lambda\gamma} = \frac{p''-c}{p'-c} > \frac{p''}{p'},$$

which is a contradiction. Therefore there can be no other symmetric equilibrium in mixed strategies.

References

Armstrong, Mark. 2008. "Interactions between Competition and Consumer Policy." *Competition Policy International* 4(1):3–53.

Cooper, Russell & Thomas Ross. 1984. "Prices, Product Qualities and Asymmetric Information: The Competitive Case." *Review of Economic Studies* 51(2):197–207.

Cruickshank, Donald. 2000. *Competition in UK Banking: A Report to the Chancellor of the Exchequer*. London: HM Treasury.

DellaVigna, Stefano. 2007. "Psychology and Economics: Evidence from the Field.". mimeo.

Ellison, Glenn. 2005. "A Model of Add-on Pricing." *Quarterly Journal of Economics* 120(2):585–637.

Financial-Services-Authority. 2006. *Levels of Financial Capability in the UK: Results of a Baseline Survey*. London: Financial Services Authority.

Gabaix, Xavier & David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *Quarterly Journal of Economics* 121(2):505–540.

Korobkin, Russell. 2003. "Bounded Rationality, Standard Form Contracts, and Unconscionability." *University of Chicago Law Review* 70:1203–1295.

Leland, Hayne. 1979. "Quacks, Lemons and Licensing: A Theory of Minimum Quality Standards." *Journal of Political Economy* 87:1328–1346.

Rosenthal, Robert. 1980. "A Model in which an Increase in the Number of Sellers Leads to a Higher Price." *Econometrica* 48(6):1575–1579.

Salop, Steven & Joseph Stiglitz. 1977. "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion." *Review of Economic Studies* 44(3):493–510.

Shleifer, Andrei. 2004. "Does Competition Destroy Ethical Behavior?" *American Economic Review* 94(2):414–418.

Varian, Hal. 1980. "A Model of Sales." *American Economic Review* 70(4):651–659.

Wolinsky, Asher. 1983. "Prices as Signals of Product Quality." *Review of Economic Studies* 50(4):647–658.