Available online at www.sciencedirect.com
ScienceDirect
Journal of Business Research 60 (2007) 858–867
JOURNAL OF BUSINESS RESEARCH

The impact of quality uncertainty without asymmetric information on market efficiency

Segismundo S. Izquierdo [a,*], Luis R. Izquierdo [b]

[a] *University of Valladolid, Spain*
[b] *University of Burgos, Spain*

Received 1 March 2006; received in revised form 1 December 2006; accepted 1 February 2007

Abstract

The market effects of quality variability and uncertainty have classically been studied in the particular context of asymmetric information, focusing on the sellers' expected behavior and the phenomenon of adverse selection. Looking instead at the consumers' expected behavior, this paper uses an agent-based model to illustrate how quality uncertainty by itself can lead to market failure, even in the absence of asymmetric information. Assuming that buyers estimate the quality of the product they buy on the basis of their experience from previous purchases, and considering quality estimation rules which are individually sensible and unbiased, this paper shows that market interaction with quality uncertainty generally produces underestimation of product quality as well as systematic drops in prices and losses of market efficiency. This study also shows that the spread of information through social networks can greatly mitigate this market failure.

Keywords: Quality uncertainty; Quality variability; Asymmetric information; Social networks

1. Introduction

Since George Akerlof's seminal paper "The Market for Lemons: Quality Uncertainty and the Market Mechanism" (Akerlof, 1970), the literature on the issue of asymmetric information and quality uncertainty has increased considerably. Following Akerlof's work, economists such as Michael Spence (Spence, 1973) and Joseph Stiglitz (Stiglitz, 2000) further developed the implications and applications of asymmetric information. They provided models that could successfully explain many otherwise surprising economic and social phenomena, such as the marked loss of market value suffered by brand-new cars on their first days of use, or the difficulties of young motorcyclists to get insurance cover, even at very high premium prices.

The works of Akerlof, Spence and Stiglitz received the Nobel Memorial Prize in Economic Sciences in 2001, and asymmetric information is now considered to be a key issue in many real markets, being one of the main paradigms underlying what is nowadays known as the economics of information (Macho-Stadler and Pérez-Castrillo, 2001; Stigler, 1961; Stiglitz, 2000).

The theory of asymmetric information has proven to be a very fruitful framework for the analysis of many types of market, but this approach does not provide a *general* answer to the original question: what is the effect of quality uncertainty in a market? The reason for this loss of generality is that, besides quality uncertainty, asymmetric information theory requires some other key assumptions that do not always necessarily hold, namely:

– There are reliable quality indicators which, before the commercial transaction takes place, are visible to only one of the potential trading partners, but not to the other (i.e. asymmetric information). For the sake of clarity, and without loss of generality, let us assume that the sellers are the possessors of privileged information.

* Corresponding author. University of Valladolid, ETS Ingenieros Industriales, Po del cauce, s/n, 47011 Valladolid, Spain.
E-mail address: segis@eis.uva.es (S.S. Izquierdo).

0148-2963/$ - see front matter © 2007 Elsevier Inc. All rights reserved.
doi:10.1016/j.jbusres.2007.02.010

– If sold at the same price, producing and selling low-quality items is more profitable than producing and selling high-quality items.
– Informed sellers present low-quality items as high-quality ones, and buyers have little or no information about the sellers' trustworthiness.
– The quality expected by every potential buyer is the market's average real quality of the product (*i.e*., perfect average information).

With these assumptions, given that uninformed buyers cannot discriminate quality before purchasing an item, one would expect high-quality and low-quality items to be sold at the same price, which would be a function of the average expected quality. Since the sales of low-quality items are more profitable at any given common price, the expectation is that low-quality items will progressively flood the market. This process would lower the average quality of the items in the market and, consequently, buyers' quality expectations and the market price.

Generally, this situation where sellers (*i.e*. the informed party) preferably offer items that are less favorable to buyers (*i.e*. the uninformed party) is known as adverse selection: it is as if the market selected adverse items for the uninformed party. For instance, a lung-illness insurance policy offered to the whole population will (unintentionally) end up selecting those individuals who are more likely to suffer from lung problems.

With adverse selection, Akerlof showed that it may even be the case that there is no possible market equilibrium at any price. Assume, for instance (Hendel and Lizzeri, 1999), that the quality q of used cars is uniformly distributed in [0, 1] and the valuation of a car of quality q is q monetary units for a potential seller and $3q / 2$ monetary units for a potential buyer; then, if the quality expected by buyers is the average quality in the market, there is no possible market equilibrium for any number of traded units but zero. To understand this, consider any equilibrium price p; the average quality of the cars offered is then $p / 2$, since only those sellers with cars of a quality below p would be willing to sell their car. In these circumstances, the buyers' valuation of a car ($3p / 4$) is lower than the price p, so no trade will take place at any possible equilibrium.

Wilson (1979, 1980) argued that markets with adverse selection might be characterized by multiple stable equilibriums. However, some years later, Rose (1993) indicated that the existence of multiple equilibriums depends critically on the distribution of quality, and that multiple equilibriums are highly unlikely for most standard probability distributions. Hendel and Lizzeri (1999) studied the interactions between new and used goods markets, and found that (in theory) the used goods market would not shut down when these interactions are considered; they then suggested that previous models overstated the distortions caused by adverse selection.

This article, similarly to Akerlof's famous case, shows that buyers' incomplete information is sufficient to cause market failure, and in some cases, even to destroy a market. In contrast to Akerlof's case, however, this study does not assume that information is necessarily asymmetric.

To illustrate this argument, this paper will analyze a model that isolates the effects of quality uncertainty from those of asymmetric information and adverse selection. Thus, to avoid confusing these different effects, the model considers that items are homogeneous at the time of sale. By assuming product homogeneity, adverse selection is necessarily avoided, since there is no *a priori* distinction between high-quality and low-quality items.

Note that product homogeneity does not mean that every item will end up providing exactly the same quality; simply that the quality distribution of every item is the same. For instance, if a product is homogeneous and an item's quality is measured by its service life, all items should have the same life expectancy. Although the literature frequently ignores this point, note that many quality features of any specific item (*e.g*. the item's service life) are random variables, since their actual value is only known when the item has been consumed. As noted by Moorthy and Hawkins (2005), products are typically consumed under noisy conditions, leading to variability in consumers' experiences. Thus, quality homogeneity at the time of purchasing should be defined in terms of quality distribution. In a practical case, product homogeneity is a valid assumption if, for instance, every item is manufactured following the same standard production process.

Related papers in the literature are those by Bergemann and Valimaki (1996), Ellison and Fudenberg (1995), and Smallwood and Conlisk (1979), who studied equilibriums in models with quality heterogeneity (brands with different quality), uncertainty and learning. Johnson and Myatt (2003) studied a Cournot model of competition in which each brand can offer multiple quality-differentiated products (quality heterogeneity, without uncertainty). Importantly, all these models focus on the effects of differences in the average quality level. In contrast, this paper isolates the effects of quality variability by assuming a constant average quality level, but different degrees of variability.

Finally, the quality expected by every potential buyer is not assumed to be the market's average real quality, since such an assumption would be difficult to hold in a number of cases. The expected quality of a product is often a subjective property, and the market's average real quality may well be unknown, or even unobservable. Even if the average quality were objective, observable, and commonly known, it is not clear that every potential buyer would use it as an unequivocal indicator to determine his or her own expected quality. In this context, this paper assumes that buyers do not form quality expectations based on the average quality of the items in the market, but based on their own past experience, and potentially influenced by the experiences of other buyers they may know. Tam (2005) explores the extent to which customers' expectations are shaped by experience in real markets.

The main, possibly striking, argument that this paper develops is that quality variability by itself can significantly damage a market if individual buyers form their quality expectations on the basis of the quality of the specific items they purchase. The paper also shows that sharing quality information through social networks can greatly reduce this damage. A simple agent-based model will illustrate these claims.

2. An agent-based model to explore the impact of quality uncertainty

This section presents a model that is a generalization of a simpler model developed by Izquierdo et al. (2005) to investigate the effects of quality uncertainty under the assumption of individual learning from personal past experience (Vriend, 2000).

This extension of the model allows buyers to learn not only from their own past experiences, but also from their social neighbors' experiences. In particular, the effect of social learning is analysed by linking buyers through a social network. In this model, the extreme case of a totally disconnected social network is equivalent to the assumption of strict individual learning from personal experiences (as in Izquierdo et al., 2005), and the extreme case of a fully connected social network is equivalent to the assumption of common knowledge of the market's average quality. This study will show that the damage caused by quality uncertainty decreases as the connectivity of the social network increases.

The following subsections explain the main features of the model. The model source code is available online at http://www.insisoc.org/research/quality, together with an applet of the model implemented in Netlogo (Wilensky, 1999), and a user guide; the reader can use the applet to replicate every experiment presented in this paper.

2.1. Supply

The supply function is constant in time. There are *num-sellers* sellers, indexed in i ($i=1,\ldots$, num-sellers). The minimum selling price for seller i is $\text{msp}_i=i$. In each trading session, every seller is allowed to sell at most one item. A seller i is willing to sell her item if the price p is no less than her minimum selling price ($p \geq \text{msp}_i$). This creates a supply function such that the number of items offered at price p ($p \geq 0$) is the integer part of p (with the additional restriction that the number of items offered cannot be greater than *num-sellers*).

2.2. Demand

The demand function in every session is formed by summing up buyers' individual reservation prices. There are *num-buyers* buyers, and the reservation price of buyer i in session n ($R_{i,n}$) is equal to her standard reservation price (SR_i) multiplied by her current expected quality ($\hat{q}_{i,n}$) for the product. Similar to the sellers, the buyers are allowed to buy at the most one item per session.

The standard reservation price SR_i for every buyer is constant throughout the simulation. Buyer i's expected quality $\hat{q}_{i,n}$ however, may vary across sessions (as detailed in Section 2.6). Each of the *num-buyers* buyers is indexed in i ($i=1, 2\ldots$ num-buyers), and buyer i has a standard reservation price SR_i equal to i. The initial expected quality $\hat{q}_{i,0}$ for every buyer is equal to 1, making every buyer's initial reservation price equal to his or her standard reservation price ($R_{i,0}=\text{SR}_i$).

Thus, given the description above, the initial demand is such that at price p ($0 < p \leq$ num-buyers), the number of items demanded is the integer part of [num-buyers$+1-p$]. Then, as trading sessions go by and buyers receive new items, they update their quality expectations and, consequently, the demand function changes.

2.3. Market design

Buyers and sellers trade in sessions. In every session, each buyer can buy one item at the most, and each seller can sell one item at the most. In every session, the market is centrally cleared at the crossing point of supply and demand. Specifically, the clearing process at any trading session n starts by sorting the buyers' individual reservation prices as follows:

$$R^1_{\bullet,n} \geq R^2_{\bullet,n} \geq \ldots \geq R^{\text{num-buyers}}_{\bullet,n}$$

Note that the upper index of the reservation prices denotes the position in the sorted list. The number of traded units in session n, v_n (for volume), is then the maximum value i such that $R^i_{\bullet,n} \geq \text{msp}_i$ and the market price p_n is taken to be:

$$p_n = 1/2[\text{Min}\,(R^{v_n}_{\bullet,n}, \text{msp}_{v_n+1}) + \text{Max}\,(R^{v_n+1}_{\bullet,n}, \text{msp}_{v_n})$$

This price-setting formula takes into account the satisfied supply and demand ($\text{msp}_{v_n} \leq p_n \leq R^{v_n}_{\bullet,n}$) and the pressure of the extramarginal supply and demand ($\text{msp}_{v_n+1} \geq p_n \geq R^{v_n+1}_{\bullet,n}$, where at least one of the inequalities is strict).

2.4. Real quality of the items

The quality q of every item follows a predetermined stationary quality distribution (*e.g.* exponential, uniform, trimmed normal). Without loss of generality the model assumes that the expected value of every distribution $E(q)$ is equal to 1.

2.5. Social network

Buyers can be connected, forming a social network. The network is created by establishing a certain number of directed links between pairs of buyers. Thus, each buyer may link to none, one, or several buyers; this (potentially empty) set of linked neighbors defines the buyer's social neighborhood. Section 3.2 investigates and discusses the relevance of the social network structure to the results.

2.6. Quality expectations updating

As outlined earlier, the initial expected quality ($\hat{q}_{i,0}$) for every buyer is equal to 1. From then onwards, in general, the buyers form their quality expectations considering both their own past experience and their social neighbors' experiences. A parameter λ_{ind} measures the sensitivity of all buyers to their own personal experiences, and a parameter λ_{soc} measures the sensitivity of all buyers to their neighbors' experiences. Thus, $\lambda_{ind} > 0$ with $\lambda_{soc}=0$ implies individual learning only.

More precisely, after every trading session n, every buyer i updates her quality expectation if and only if:

- she has bought an item and she somewhat considers her own experience ($\lambda_{ind} > 0$), or
- someone in her social neighborhood has bought an item and she somewhat considers her neighbors' experiences ($\lambda_{soc} > 0$).

When buyer i updates her expectations, she does so according to the following rules:

a) If both buyers i and someone in her neighborhood has purchased an item:

$$\hat{q}_{i,n+1} = \hat{q}_{i,n} + \lambda_{ind} \cdot (q_{i,n} - \hat{q}_{i,n}) + \lambda_{soc} \cdot (\overline{q}_{i,n} - \hat{q}_{i,n})$$

where $q_{i,n}$ is the quality of the item received by buyer i in session n, $\overline{q}_{i,n}$ is the average quality of the items received by the buyers in i's social neighborhood, and λ_{ind} and λ_{soc} are the individual and social learning rate respectively. Note that the learning rates measure the responsiveness of the buyers' quality estimates to new data.

b) If buyer i has purchased an item, but none in her neighborhood has, then:

$$\hat{q}_{i,n+1} = \hat{q}_{i,n} + \lambda_{ind} \cdot (q_{i,n} - \hat{q}_{i,n})$$

c) If buyer i has not purchased an item, but someone in her neighborhood has, then:

$$\hat{q}_{i,n+1} = \hat{q}_{i,n} + \lambda_{soc} \cdot (\overline{q}_{i,n} - \hat{q}_{i,n})$$

Values in the range $0 \leq \lambda_{ind}, \lambda_{soc} \leq 1$ are considered, but note that combinations of values such that $(\lambda_{ind}+\lambda_{soc}) > 1$ could mean an over-reaction of buyers to new quality data. The Appendix Section discusses the different interpretations of this additive learning model.

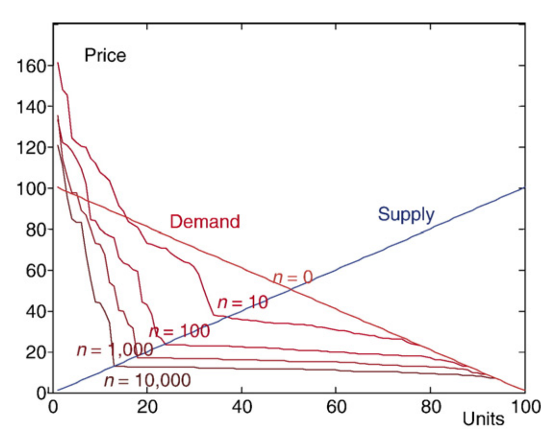

Fig. 1. Effects of quality variability on demand. Quality distribution: $q \sim U[0, 2]$. There are 100 unconnected buyers (individual learning). The initial demand ($n=0$) is linear.

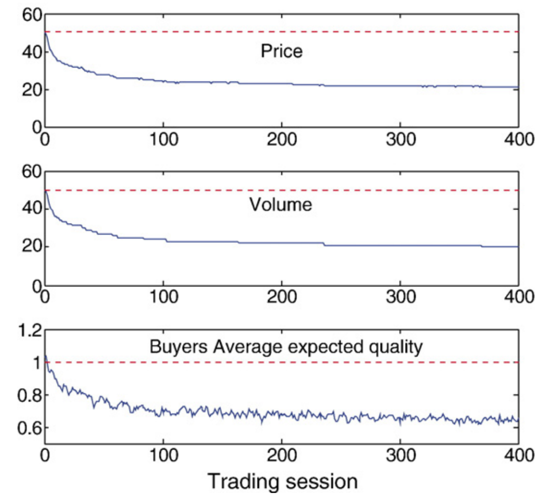

Fig. 2. Effects of quality variability on price level (top), traded volume (middle) and average expected quality (bottom). The dotted line shows the reference situation (no quality variability).

3. Results: market failure

3.1. Individual learning

This section discusses the individual learning case ($\lambda_{soc}=0$), which is based on a model developed by Izquierdo et al. (2005), and provides two propositions about the dynamics of the individual-learning models.

With individual learning, buyers update their expected quality only when they (individually) receive a new item and observe its quality. In each session, the market is centrally cleared at the crossing point of supply and demand, and all the buyers who have bought an item update their quality expectations according to their experience with the item. A key assumption of this model is that those buyers who do not get items do not update their quality expectations: new information about the product is only acquired by new purchases.

Note that, in these conditions, if there were no quality variability, the initial market equilibrium would last indefinitely.

As a particular case of a market with individual learning, consider an initial situation ($n=0$) such as the one shown in Fig. 1, which corresponds to a parameterization with 100 buyers and 100 sellers where the quality q of every item follows a uniform quality distribution $q \sim U[0, 2]$. The reference conditions (*i.e.* no quality variability) are price=50.5, traded volume=50. These conditions would last indefinitely maintained if there were no product variability. However, in this model there is quality variability and individual quality learning.

Surprisingly, in this model with symmetric quality variability and unbiased learning rules, inefficient market dynamics emerge, prices drop below reference conditions, and buyers systematically underestimate the actual quality of the product.

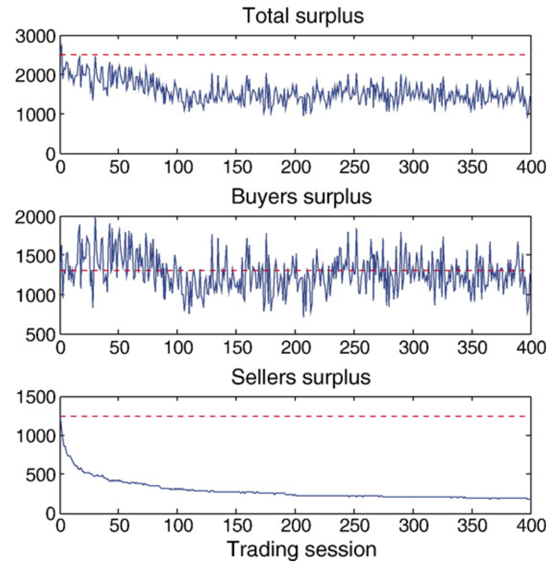

Fig. 3. Effects of quality variability on total surplus (top), buyers' surplus (middle) and sellers' surplus (bottom). The dotted line shows the reference situation (no quality variability).

Fig. 1 shows results corresponding to a learning rate $\lambda_{\text{ind}}=0.5$. The degeneration of the demand function is clearly visible from the early periods. After a certain number of periods the demand function seems rather stable and the results of consecutive trading sessions look very similar. However, as this paper will show later, with these conditions and given enough time, no trading would eventually take place.

Figs. 2 and 3 show a pattern of decreasing prices, decreasing expected quality, monotonously decreasing number of traded units, and loss of efficiency which is consistent throughout simulations for different numbers of players (100 buyers and 100 sellers in the figures), for different values of λ_{ind} (0.5 in the figures) and for different quality distributions ($U[0, 2]$ in the figures). This is proven mathematically below.

In this simulated market, because of the drop in sales and prices, there can be a great loss of surplus, especially for sellers (Fig. 3). The seller's surplus in a transaction between a seller and a buyer is the difference between the price of the item sold (seller's income) and the seller's minimum selling price for that item (this is the minimum price that the seller would be willing to accept in exchange for the item, which is usually the item's marginal cost of production); the buyer's surplus is the difference between the maximum price that the buyer is willing to pay for the item (reservation price, or marginal value) and the price actually paid (cost).

Note that in this model the average quality of the items is constant ($E(q)=1$) and the buyers' quality learning rule is unbiased but, as trading sessions go by, most buyers perceive a quality lower than the real one, and the average perceived quality is consistently lower than the real average quality.

With individual learning, the market price provides a dynamic threshold that separates the buyers who receive an item and update their quality expectations from the buyers who do not update their quality expectations. The lower a buyer's quality expectations are, the less likely the chance that she will buy a new item and update her expectations, so low-quality expectations are more likely to be maintained than high-quality ones. For every buyer, the dynamics of quality expectations are conditioned on the expected-quality value, and the lower this value gets, the less likely it is to evolve.

The essence of the phenomenon is more clearly understood when assuming that supply is horizontal at a given price level X (any amount of items can be sold at price X, but not below). If by purchasing a series of bad items a buyer's reservation price can drop below X, she will stop buying the product for good.

More generally than these particular cases, consider any market model M such that:

- Buyers and sellers trade in sessions. In each session, each buyer can buy one item at the most. No item sells at a price lower than its minimum selling price or higher than its buyer's reservation price.
- Buyers' reservation prices depend on their current quality expectations for the product. Buyers who do not get a new item do not update their quality expectations (or their reservation price).
- The supply function (the number of items whose minimum selling price is lower than any given price p) is constant in time (*i.e.* supply does not change over trading sessions).
- The market clearing mechanism is such that a common price is set where supply and demand intersect, leaving no buyer or seller unsatisfied (*i.e.* every buyer with a reservation price higher than the market price receives an item, and every item with a minimum selling price lower than the market price is sold).

Then, starting from any initial conditions, if quality variability is introduced, the following two propositions hold (the Appendix Section provides proofs).

Proposition 1. *The number of traded units in a market model M is monotonously decreasing in time.*

Note that Proposition 1 holds for any learning rule and any quality distribution. The main result of Proposition 1 summarizes as follows: if supply is constant and those buyers who do not purchase an item do not change their reservation prices, then, starting from any initial situation, the number of tradable units is monotonously decreasing. Whether in the long-term the market will totally collapse or whether the market will reach a stable equilibrium depends on the quality distribution and on the particular learning rules used by the buyers.

Proposition 2. *Let $Hmsp_n$ be the highest minimum selling price of all the traded units at session n in a market model M. If at every trading session n there is a positive (bounded away from 0) probability that some reservation price(s) will (in a finite number of sessions) drop below $Hmsp_n$, then the market will eventually collapse.*

In particular, consider the model that Figs. 1–3 show. Given the quality distribution $q \sim U[0, 2]$ and the quality expectations updating rule, there is a positive probability for any buyer's reservation price to fall bellow Hmsp_n in every session (the minimum value for Hmsp_n is 1), so this market will eventually collapse.

3.2. Social learning

The assumption that buyers' expected quality is only based on their own past experience may not seem realistic in those markets in which information can be easily shared among consumers, or in which there is reliable aggregate information on the product's quality available to the general public (*e.g.* journals, magazines, public reports or discussion forums). In particular, information about many products is now easily accessible on the Internet, influencing consumers' shopping behavior (Senecal et al., 2005).

The market damage caused by quality uncertainty with individual learning is due to the fact that a buyer obtains new information about a product's quality only after a new purchase. As lower quality expectations imply lower chances of purchasing a new item, long-sustained low-quality expectations are favored over long-sustained high-quality expectations (assuming that the learning rule is not biased). The dynamics of quality expectations are asymmetric, because they are conditioned on the value of the expected quality, and the lower this value gets, the less likely it is to change.

In a context of shared information, assuming buyers' responsiveness to new data remains approximately constant, one can expect two combined effects: first, less variability on every buyer's quality estimates over time, as the estimates would be based on more data, and second, the flow of information obtained by every buyer would be less conditioned by their own reservation price, as they could be receiving new information even when they (individually) did not purchase a new item. As a consequence of both effects, one would expect a lower damage caused by quality uncertainty.

The rest of this section provides some simulation results from a model of information sharing through randomly generated social networks, where links are created between randomly selected pairs of buyers. Section 3.3 discusses the robustness of the results to changes in the network-generating algorithm.

The extreme case of a fully connected social network would be equivalent to the assumption of common knowledge of the market's average quality. The damage caused by quality uncertainty with individual learning usually decreases considerably as the connectivity of the social network grows (Fig. 4 shows representative runs). The general pattern is the same for different quality distributions: uniform, trimmed normal, or exponential (as in the following figures).

Fig. 5 shows the average number of traded units (sales) at session 500 across 1000 random networks for various combinations of number of links (network connectivity) and individual learning rate λ_{ind}. The variability across runs is limited (the standard deviation for sales is less than 3.6 units in every case; the standard error for the average values shown in

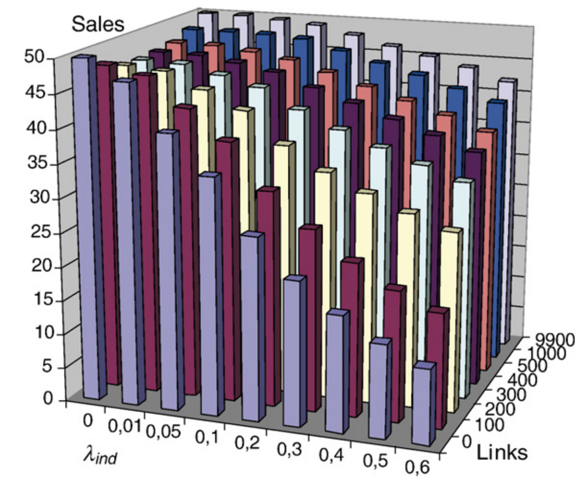

Fig. 5. Average (across 1000 random networks in every case) sales at trading session 500, measured in models with different λ_{ind} and number of random links, with 100 buyers, 100 sellers, $\lambda_{\mathrm{soc}}=0.4$ and $q \sim \exp(1)$.

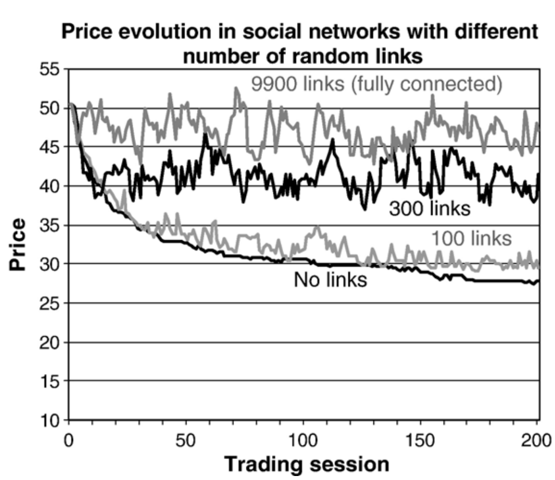

Fig. 4. Price evolution in 4 random social networks with 100 buyers, 100 sellers, and different number of random links. Quality distribution $q \sim \exp(1)$, $\lambda_{\mathrm{ind}}=0.25$, $\lambda_{\mathrm{soc}}=0.25$.

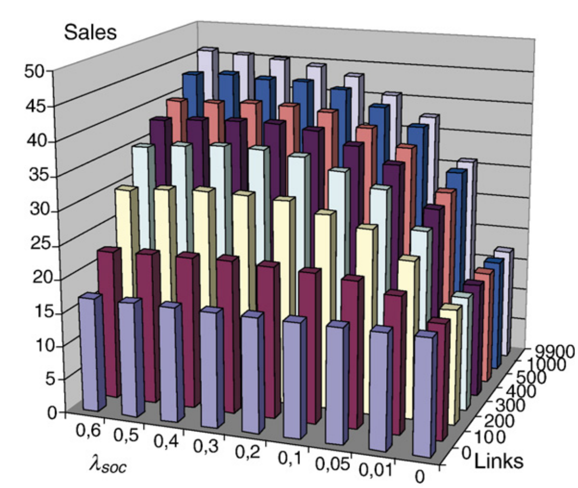

Fig. 6. Average (across 1000 random networks in every case) sales at trading session 500, measured in models with different λ_{soc} and number of random links, with 100 buyers, 100 sellers, $\lambda_{\mathrm{ind}}=0.4$ and $q \sim \exp(1)$.

the graph is less than 0.12 in every case). Similar patterns show on average expected qualities, prices and sellers' surplus.

Note that for higher λ_{ind} the quality estimations are more variable, causing an effect similar to that of a higher quality variability (*i.e.* lower prices and number of sales, more marked quality underestimation, and higher losses of market efficiency). Note also that, as the number of links in the social network grows (shared information), and more than one quality experience is considered when updating the expectations, the damaging effects of quality variability can be greatly reduced (Figs. 5 and 6): sharing information usually reduces the variability of quality expectations, and it also reduces the dependence of the flow of new information on the value of the individual expected quality. Fig. 6 shows the (non-linear) effect of λ_{soc} keeping the value of λ_{ind} constant.

3.3. Robustness with respect to different network structures

Randomly generated networks can be a good way of testing the robustness of a market effect with respect to changes in the network structure (after all, given a certain number of links, the random procedure used can generate any possible network design). However, different algorithms for network creation will lead to different statistical regularities in the behavior of the resulting networks.

The validity of the results is tested using some other network-generating algorithms, such as the preferential attachment rule of Barabási and Albert (1999) as described by Newman (2003, section VII B). The general results presented in this paper are robust with respect to changes in the network-generating algorithm, but, however, the same network-generating algorithm can give rise to particular networks with very different behaviors. For instance, consider a star-shaped network-generating algorithm such that one buyer is randomly selected to be the centre of the star and a bidirectional link is created between this buyer and each one of the other buyers. The properties of the market in a star-shaped network critically depend on the behavior of the buyer in its center. If the central buyer is a frequent consumer, all the other buyers will be updating their quality expectations frequently through her, and the market will not suffer much from the long-lasting loss of confidence effect. However, if the central buyer only purchases an item occasionally, she will only update the market expectations occasionally, between periods of increasing loss of confidence.

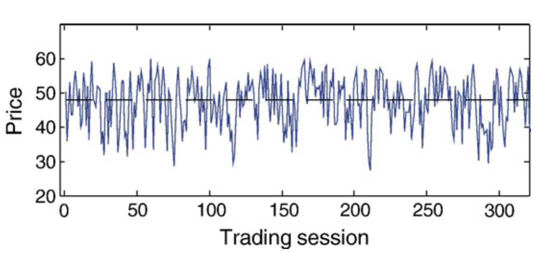

Fig. 7. Price evolution in a market model with a star-shaped social network. The standard reservation value of the central buyer is 63. Conditions: 100 buyers, 100 sellers, $\lambda_{ind}=0.4$, $\lambda_{soc}=0.4$, $q\sim U(0, 2)$.

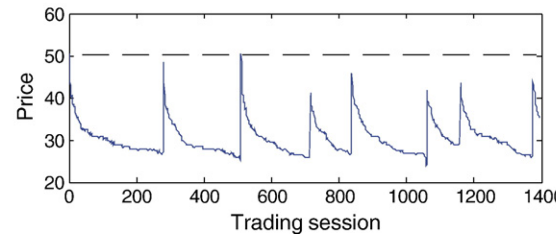

Fig. 8. Price evolution in a market model with a star-shaped social network. The standard reservation value of the central buyer is 25. Conditions: 100 buyers, 100 sellers, $\lambda_{ind}=0.4$, $\lambda_{soc}=0.4$, $q\sim U(0, 2)$.

To illustrate this last point, consider a market with 100 buyers (standard reservation prices = 1, 2..., 100) and 100 sellers (minimum selling prices = 1, 2..., 100). The reference conditions (no quality variability) for the price and sales in this market are close to 50. Fig. 7 shows the evolution of the prices (with quality variability) in a star-shaped social network whose central buyer has a standard reservation price of 63.

Now consider the same sellers and buyers, also embedded in a star-shaped social network, but the central buyer now has a standard reservation price of 25. As before, the reference conditions (no quality variability) for the price and sales are close to 50, but in general, the central buyer will not purchase an item unless the price drops close to 25. Fig. 8 depicts the evolution of the prices in one of these networks, with periods of loss of confidence in between shocks caused by purchases of the central buyer. The shocks are usually upward because the quality of the new items of the central buyer is usually above the (depressed) average expectations. After a price recovery, the central buyer will stop buying until the prices decrease to the level of her reservation price. Thus, this example shows that one single (stochastic) network-generating algorithm can lead to specific networks displaying dramatically different behavior.

4. Discussion and conclusions

This paper analyzes the impact of quality variability on markets. This analysis has classically been carried out in the particular framework of asymmetric information and adverse selection. Although extremely useful, this framework requires two important assumptions (asymmetric information and buyers' quality expectations equal to the average market quality, *i.e.* common knowledge of the market's real average quality), which do not necessarily hold in every case of quality uncertainty. Besides, the effect of quality uncertainty by itself in the asymmetric information model is difficult to isolate from the effects of the other particular assumptions of that model. This paper investigated the effect of quality uncertainty in a more general framework where information is not necessarily asymmetric and buyers estimate product quality on the basis of past experiences.

Considering this framework, the model recognizes that quality expectations may not be common to every buyer, but that they may rather depend on buyers' personal experiences with the product. The assumption of one single homogeneous quality distribution for every item is not instrumental in order to observe the effects of quality variability discussed here, but this assumption ensures a neat distinction between the effects of quality variability in general and its effects in the particular case where there is also adverse selection.

The striking fact this paper illustrates is that quality variability combined with the assumption that buyers estimate product quality on the basis of their past experience can significantly damage the market, especially so when quality variability is high and quality information is not widely spread. This effect is not due to buyers' risk aversion (which has not been included in the model), but to a generally sustained underestimation of the product quality.

The underlying reason for this phenomenon is that buyers who happen to receive a low-quality item are less likely to buy new items of that particular kind – and consequently less likely to update their low-quality perception of the product – than buyers who receive a higher quality item. Thus, low-quality expectations tend to persist longer than high quality ones. New purchases provide new information about product quality, but new purchases are conducted primarily by buyers who have higher quality expectations.

The extreme case of no information sharing plus high-quality variability can completely destroy a market, but this paper also shows that making aggregate information available, or sharing information through a social network, can greatly mitigate these damaging market effects. When information is shared, buyers with low expectations may still be able to revise them through their social links.

The model of quality uncertainty discussed in this paper could be extended to include other features such as buyers' risk aversion or asymmetric spreading of bad and good news in social networks. Note, however, that the main point this model demonstrates is precisely that these other features are not necessary for quality uncertainty to damage a market and to undermine the confidence in the product.

From a practical point of view, when analyzing a market, the aggregate results of the loss-of-confidence effects discussed here may be difficult to distinguish from the effects of adverse selection in many cases, but specific market characteristics can assist in assessing the relative importance of each effect. For instance:

- Adverse selection will rarely be an issue if the quality differences among suppliers are not large (e.g. in commodity markets, monopolies, or industries with standard processes). In this case, quality variability could still damage the market because of the effect described in this paper or because of buyers' risk aversion.
- The case for adverse selection is also weak when there are few agents in the market and frequent interactions among them, because of the role of reputation (see the discussion in Kirman and Vriend, 2001 for the wholesale fish market in Marseille).
- The validity of the assumption of common knowledge of average quality is likely to depend on the number and the frequency of individual purchases. It may be easy to calculate the average quality in markets where individual buyers can check the quality of a large number of items (e.g., insurance companies), but in other markets this calculation may be more difficult (e.g., used cars markets).
- In a given market, the importance of personal past experiences (as apposed to aggregate indicators) in people's purchasing behavior can be empirically tested, either through surveys or through controlled experiments.
- The explanation that this article puts forward predicts the average expected quality to be lower than the real average quality, but the model based on asymmetric information assumes that every buyer knows the real average quality. Thus, detecting such a difference between real and perceived quality would be an indication of the potential presence of the effect investigated here.

Finally, this article discusses a loss-of-confidence effect due to quality variability at industry level while assuming product homogeneity. A somewhat related situation is that of different firms who provide items with a similar average quality but with a different quality variability. In this situation the loss-of-confidence effect due to quality variability can be critical for individual companies, and some common marketing policies can be justified under this perspective. For instance, it is sometimes observed in the food market that some retailers provide warranties that reimburse the cost of any defective item and replace the item with a new one. The rationale behind this policy is not only reassuring the buyers' *a priori* confidence in the product's quality (a cheap good warranty is a clear signal of good quality), but also restoring the buyer's approval of the product if she happens to receive a defective item, and prevent her from switching to another brand. It is worth noting, however, that there are some markets where even satisfied customers switch between brands very frequently, as discussed by Arnold et al. (2005) and by Chiu et al. (2005).

Acknowledgments

The authors gratefully acknowledge financial support from the Scottish Executive Environment and Rural Affairs Department and from the SocSimNet project 2004-LV/04/B/F/PP. The authors would also like to thank Dale Rothman, Cesáreo Hernández and the participants in the 2006 workshop Agent Based Models of Market Dynamics and Consumer Behavior at the University of Surrey for their useful comments.

Appendix

A.1. Proof of Proposition 1

Let $[\mathrm{msp}^1, \mathrm{msp}^2, \ldots]$ be the vector of the minimum selling prices of the items in the market, sorted out in ascending order, and let $\left[R^1_{,n}, R^2_{,n}, \ldots, R^{\mathrm{num-buyers}}_{,n}\right]$ be the vector of the *num-buyers* reservation prices (one for each buyer) at session n, sorted out in descending order.

Let v_n be the number of traded units at session n. Given the definition of v_n, for any number of units i with $0 < i <$ num-buyers, the following holds:

$$R^{i}_{\bullet\bullet,n} < \mathrm{msp}^{i} \Leftrightarrow v_n < i \qquad (1)$$

In particular, for $i = v_n + 1$,

$$R^{v_n+1}_{\bullet\bullet,n} < \mathrm{msp}^{v_n+1}. \qquad (2)$$

At session n there are (num-buyers $-v_n$) buyers who do not purchase any item and whose reservation prices are not higher than $R^{v_n+1}_{\bullet\bullet,n}$. As those buyers will not change their reservation price for the next session, $R^{v_n+1}_{\bullet\bullet,n+1} \le R^{v_n+1}_{\bullet\bullet,n}$, from where, using Eq. (2), $R^{v_n+1}_{\bullet\bullet,n+1} < \mathrm{msp}^{v_n+1}$, and using Eq. (1), $v_{n+1} < v_n + 1$, which implies $v_{n+1} \le v_n$. □

A.2. Proof of Proposition 2

Let us call purchasers those buyers who acquire an item in a particular session. Given the market clearing mechanism, the items that are actually exchanged in session n are the v_n items with a lower minimum selling price, and therefore:

$$\mathrm{Hmsp}_n = \mathrm{msp}^{v_n}.$$

Now divide the set of buyers in a given session n into two subgroups:

- The Potential purchasers subgroup: those buyers with reservation prices higher than or equal to msp^{v_n}. Each purchaser in a session has to be in this subgroup.
- The Outsiders subgroup: those buyers with reservation prices lower than msp^{v_n}. Nobody in this subgroup can be a purchaser, and therefore nobody in this group will update her reservation price.

The following will prove that, given any situation where some degree of trade takes place, the number of units decreases with probability 1 (not necessarily in the following session, but eventually). First, note that the number of traded units cannot increase, as demonstrated in Proposition 1. Note also that while the number of traded units remains equal to v_n, the highest minimum selling price remains equal to msp^{v_n}. Therefore, unless the number of traded units decreases, the highest minimum selling price will remain equal to msp^{v_n}. This means that while the number of traded units remains equal to v_n, the individuals in the Outsiders group will not be able to purchase any item, and will therefore stay in this group. On the other hand, the individuals in the Potential purchasers group may move to the Outsiders group, and this will happen with probability 1, since, by assumption, in every session $m \ge n$ there is a positive (bounded away from 0) probability that some reservation price (s) will (in a finite number of sessions) drop below $\mathrm{Hmsp}_m = \mathrm{msp}^{v_m} = \mathrm{msp}^{v_n}$. When the number of individuals in the potential purchasers group drops below v_n, the number of traded units will necessarily decrease.

A necessary condition for total collapse is that at some point every buyer's reservation price has a positive probability of falling below the minimum possible selling price. □

A,3. Note on the quality expectations updating rule of Section 2.6

The additive model of Section 2.6:

$$\hat{q}_{i,n+1} = \hat{q}_{i,n} + \lambda_{\mathrm{ind}} \cdot (q_{i,n} - \hat{q}_{i,n}) + \lambda_{\mathrm{soc}} \cdot (\overline{q}_{i,n} - \hat{q}_{i,n})$$

is equivalent to a model in which the quality updating factor is a linear combination of the social and the individual observations:

$$\hat{q}_{i,n+1} = \hat{q}_{i,n} + \alpha_1 \cdot [(\alpha_2 \cdot q_{i,n} + (1-\alpha_2) \cdot \overline{q}_{i,n}) - \hat{q}_{i,n}]$$

$$\lambda_{\mathrm{ind}} = \alpha_1 \cdot \alpha_2 \quad \lambda_{\mathrm{soc}} = \alpha_1 \cdot (1-\alpha_2).$$

This model is also equivalent to one in which the individual (social) observation modifies the expected quality first, after which the social (individual) observation modifies the new expectation:

$$\hat{q}_{i,n+1} = \hat{q}_{i,n} + \alpha_{\mathrm{ind}} \cdot (q_{i,n} - \hat{q}_{i,n}) + \alpha_{\mathrm{soc}} \cdot [\overline{q}_{i,n} - (\hat{q}_{i,n} + \alpha_{\mathrm{ind}} \cdot (q_{i,n} - \hat{q}_{i,n}))$$

$$\lambda_{\mathrm{ind}} = \alpha_{\mathrm{ind}} \cdot (1-\alpha_{\mathrm{soc}}) \quad \lambda_{\mathrm{soc}} = \alpha_{\mathrm{soc}}.$$

References

Akerlof GA. The market for lemons: quality uncertainty and the market mechanism. Q J Econ 1970;84:488–500.

Arnold MJ, Reynolds KE, Ponder N, Lueg JE. Customer delight in a retail context: investigating delightful and terrible shopping experiences. J Bus Res 2005;58(8):1132–45.

Barabási AL, Albert R. Emergence of scaling in random networks. Science 1999;286:509–12.

Bergemann D, Valimaki J. Learning and strategic pricing. Econometrica 1996;64(5):1125–49.

Chiu HC, Hsieh YC, Li YC, Lee M. Relationship marketing and consumer switching behaviour. J Bus Res 2005;58(12):1681–9.

Ellison G, Fudenberg D. Word of mouth communication and social learning. Q J Econ 1995;110(1):93–125.

Hendel I, Lizzeri A. Adverse selection in durable goods markets. Am Econ Rev 1999;89(5):1097–115.

Izquierdo SS, Izquierdo LR, Galán JM, Hernández C. Market failure caused by quality uncertainty. In: Mathieu P, Beaufils B, Brandouy O, editors. Artificial economics — lecture notes in economics and mathematical systems 564. Berlin: Springer-Verlag; 2005.

Johnson JP, Myatt DP. Multiproduct quality competition: fighting brands and product line pruning. Am Econ Rev 2003;93(3):748–74.

Kirman AP, Vriend NJ. Evolving market structure: an ACE model of price dispersion and loyalty. J Econ Dyn Control 2001;25(3–4):459–502.

Macho-Stadler I, Pérez-Castrillo JD. An introduction to the economics of information. incentives and contracts (2nd ed.). Oxford University Press; 2001.

Moorthy S, Hawkins SA. Advertising repetition and quality perception. J Bus Res 2005;58(3):354–60.

Newman MEJ. The structure and function of complex networks. SIAM Rev 2003;45:167–256.

Rose C. Equilibrium and adverse selection. Rand J Econ 1993;24(4):559–69.

Senecal S, Kalczynski PJ, Nantel J. Consumers' decision-making process and their online shopping behavior: a clickstream analysis. J Bus Res 2005;58 (11):1599–608.

Smallwood DE, Conlisk J. Product quality in markets where consumers are imperfectly informed. Q J Econ 1979;93(1):1–23.

Spence AM. Job market signaling. Q J Econ 1973;83:355–77.

Stigler GJ. The economics of information. J Polit Econ 1961;69:213–25.

Stiglitz JE. The contributions of the economics of information to twentieth century economics. Q J Econ 2000;115(4):1441–78.

Tam JLM. Examining the dynamics of consumer expectations in a Chinese context. J Bus Res 2005;58(6):777–86.

Vriend N. An illustration of the essential difference between individual and social. learning, and its consequence for computational analyses. J Econ Dyn Control 2000;24:1–19.

Wilensky U. NetLogo. http://ccl.northwestern.edu/netlogo/. Center for Connected Learning and Computer-Based Modeling. Evanston, IL: Northwestern University; 1999.

Wilson CA. Equilibrium and adverse selection. Am Econ Rev 1979; 69:313–7.

Wilson CA. The nature of equilibrium in markets with adverse selection. Bell J Econ 1980;11:108–30.