# The impact on market efficiency of quality uncertainty without asymmetric information

Segismundo S. Izquierdo, Luis R. Izquierdo

Segismundo S. Izquierdo (corresponding author)
University of Valladolid
ETS Ingenieros Industriales
Pº del cauce, s/n
47011Valladolid
segis@eis.uva.es
Phone: 00 34 983 183873 / 00 34 617 373043

Luis R. Izquierdo
The Macaulay Institute
Craigiebuckler,
Aberdeen AB15 8QH
l.izquierdo@macaulay.ac.uk

**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 behaviour and the phenomenon of adverse selection. Looking instead at the consumers' expected behaviour, we use 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 using their past experience from previous purchases, and considering quality estimation rules which are individually "sensible" and unbiased, market interaction is shown to produce general underestimation of product quality, as well as systematic drops in prices and losses in market efficiency. It is also shown that the spread of information through social networks can greatly mitigate this market failure.

**Keywords**: quality uncertainty, quality variability, asymmetric information, social networks

# The impact on market efficiency of quality uncertainty without asymmetric information

## 1 Introduction

Since George Akerlof's seminal paper "The Market for Lemons: Quality Uncertainty and the Market Mechanism" (Akerlof 1970), there has been a growing literature on the issue of asymmetric information and quality uncertainty. 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, providing 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 were awarded 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" (Stigler 1961, Stiglitz 2000, Macho-Stadler and Pérez-Castrillo 2001).

The theory of asymmetric information has proven to be a very fruitful framework for the analysis of many markets, but it 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, the asymmetric information theory requires some other key assumptions which do not always necessarily hold, mainly:

- There are reliable quality indicators which, before the commercial transaction has been made, can be observed by only one of the potential trading partners, but not by the other (*i.e.* asymmetric information). For the sake of clarity, and without loss of generality, let us assume that it is the sellers who have privileged information.

- If sold at the same price, producing and selling low-quality items is more profitable than producing and selling high-quality ones.

- Informed sellers present low-quality items as high-quality ones, and buyers have little or no information about 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 can not 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 sales of low-quality items are more profitable at any given common price, it is expected 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 those items that are less favourable 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.

As shown by Akerlof, when there is adverse selection 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 offered cars is then $p/2$, since only those sellers with cars of quality below $p$ would be willing to sell their car. In these conditions, buyers' valuation of a car ($3p/4$) is lower than the price $p$, so there is no trade at any possible equilibrium.

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

In this paper, similarly to Akerlof's famous case, we show that buyers' incomplete information is sufficient to cause market failure, and even destroy a market in some cases. In contrast to Akerlof's case, however, we do not assume that information is necessarily asymmetric.

To illustrate our argument, we analyse a model in which the effects of quality uncertainty are isolated from those of asymmetric information and adverse selection. Thus, to avoid confusing these different effects, we consider that items are homogeneous at the time of being sold. 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 by "product homogeneity" we do not mean that every item will end up providing exactly the same quality; we mean that the quality distribution of every item is just 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 expected life). Though this point is frequently ignored, 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. Thus, quality homogeneity at the time of purchasing should be defined in terms of quality distribution. In a practical case, we would recognise product homogeneity as a valid assumption if, for instance, every item is manufactured following the same standard production process.

Related papers in the literature are those by Smallwood and Conlisk (1979), Ellison and Fudenberg (1995), and Bergemann and Valimaki (1996), who study equilibria 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, in all these models products can be offered at different average levels of quality. In contrast to this, in this paper we isolate the effects of quality variability by assuming a constant average level of quality.

Finally, we do not wish to assume that the quality expected by every potential buyer is the market's average real quality, since we find such an assumption 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 their own expected quality. In this context, the assumption in our model is 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.

The main, possibly striking argument that we show and develop in this paper is that quality variability by itself can significantly damage a market if individual buyers form their quality expectations based on the quality of the specific items they purchase. We also show that sharing quality information through social networks can greatly reduce this damage. These points will be illustrated using a simple agent-based model.

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

In this section we present our model, which is a generalisation 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).

In our extension, we allow buyers to learn not only from their own past experience, but also from their social neighbours' experiences. In particular, we analyse the effect of social learning 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. The damage caused by quality uncertainty will be shown to decrease as the connectivity of the social network increases.

The following subsections explain the main features of our 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 that we present in this paper.

## 2.1 Supply

The supply function is constant in time. There are *num-sellers* sellers indexed in $i$ ($i$ = 1,..., *num-sellers*) with minimum selling price for seller $i$ being $msp_i = i$. In each trading session every seller may 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 ($p \geq msp_i$). 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. As with sellers, buyers may buy at most one item per session.

The standard reservation price $SR_i$ for every buyer is constant throughout the simulation. Her 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 ... *num-buyers*), and buyer $i$ has 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 their standard reservation price ($R_{i,0} = SR_i$).

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 each session, each buyer can buy at most one item, and each seller can sell at most one item. 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 buyers' individual reservation prices as follows:

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

Note that the upper index in 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 $R^i_{\bullet,n} \geq msp_i$ , and the market price $p_n$ is taken to be:

$$p_n = \tfrac{1}{2}\left[\mathrm{Min}(R^{v_n}_{\bullet,n}, msp_{v_n+1}) + \mathrm{Max}(R^{v_n+1}_{\bullet,n}, msp_{v_n})\right]$$

This price-setting formula takes into account the satisfied supply and demand ($msp_{v_n} \leq p_n \leq R^{v_n}_{\bullet,n}$) and the pressure of the extramarginal supply and demand ($msp_{v_n+1} \geq p_n \geq R^{v_n+1}_{\bullet,n}$ , 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 we assume 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 neighbours defines the buyer's social neighbourhood.

## 2.6 Quality expectations updating

As mentioned before, the initial expected quality ($\hat{q}_{i,0}$) for every buyer is equal to 1. From then onwards, in general, buyers form their quality expectations considering both their own past experience and their social neighbours' experiences. A parameter $\lambda_{ind}$ measures the sensitivity of all buyers to their own personal experiences, and a parameter $\lambda_{soc}$ measures the sensitivity of all buyers to their neighbours' 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 neighbourhood has bought an item and she somewhat considers her neighbours' experiences ($\lambda_{soc} > 0$).

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

a) If both buyer *i* and someone in her neighbourhood 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 buyers in *i*'s social neighbourhood, and $\lambda_{ind}$ and $\lambda_{soc}$ are the individual and social "learning rate" respectively. Note that the learning rates measure the responsiveness of buyers' quality estimates to new data.

b) If buyer *i* has purchased an item but none in her neighbourhood has:

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

c) If buyer *i* has not purchased an item but someone in her neighbourhood has:

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

We consider values in the range $0 \le \lambda_{ind}, \lambda_{soc} \le 1$, but note that combinations of values such that $(\lambda_{ind} + \lambda_{soc}) > 1$ could mean "over-reaction" of buyers to new quality data. Different interpretations of this additive learning model are discussed in the appendix.

# 3 Results: Market failure

## 3.1 Individual learning

We begin by discussing the individual learning case ($\lambda_{soc} = 0$), which is based on a model developed by Izquierdo et al. (2005). Izquierdo et al. (2005) described the market dynamics in their model and tested the robustness of their results to various market mechanisms. In this section we show the main results for our model, and we prove two propositions about the dynamics of these 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 just bought. 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, there would be a sustained market equilibrium at the crossing point of supply and demand which would be preserved indefinitely.

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

Surprisingly, in our 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.

Figure 1 shows some results corresponding to a learning rate $\lambda_{ind} = 0.5$. The degeneration of the demand function can be clearly seen 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. We show later, however, that with these conditions and given enough time, no trading would eventually take place.

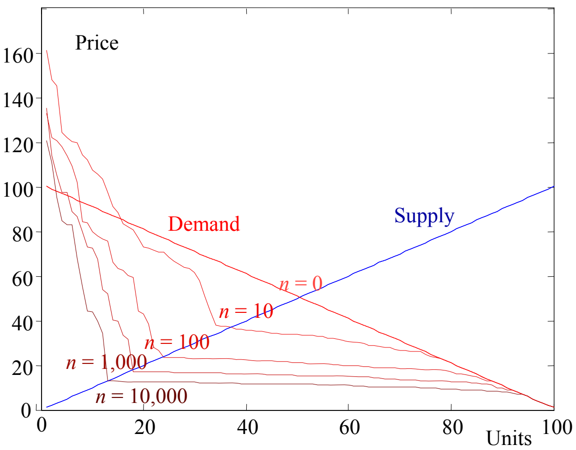


Figure 1. Effects of quality variability on demand. Quality distribution: $q \sim U[0, 2]$. There are 100 unconnected buyers (individual learning). The initial demand ($n = 0$) is linear.

The general pattern shown in Figure 2 and Figure 3 (decreasing prices, decreasing expected quality, monotonously decreasing number of traded units, and loss of efficiency) is consistent throughout simulations for different numbers of players (100 buyers and 100 sellers in the figures), for different values of $\lambda_{ind}$ (0.5 in the figures) and for different quality distributions (U[0, 2] in the figures). We prove this mathematically below.

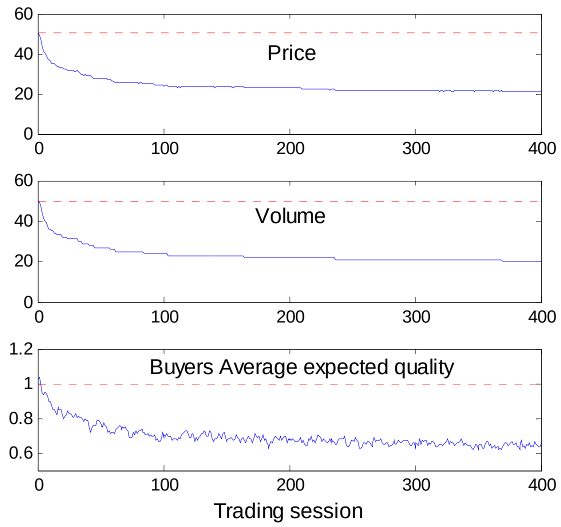


Figure 2. Effects of quality variability on price level (top), traded volume (middle) and average expected quality (bottom). The dotted line shows the reference situation (no quality variability).

In our simulated market, because of the drop in sales and prices, there can be a great loss of surplus, especially for sellers (Figure 3). The seller's surplus in a transaction between a seller and a buyer is the difference between the price of the sold item (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, and it is usually the item's marginal cost, if the item is to be produced); the buyer's surplus is the difference between the maximum price that the buyer would have paid for the item (reservation price, or marginal value) and the price actually paid (cost).

Note that in our 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.

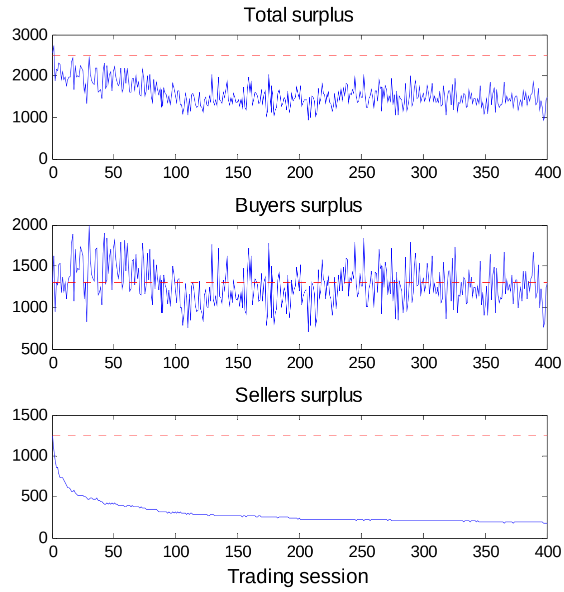


Figure 3. Effects of quality variability on total surplus (top), buyers' surplus (middle) and sellers' surplus (bottom). The dotted line shows the reference situation (no quality variability).

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

The essence of the phenomenon is more clearly understood if we assume 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 at most one item. No item is sold at a price lower than its minimum selling price or greater 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 (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 reservation price greater than the market price is given an item, and every item with minimum selling price lower than the market price is sold)

Then, for any initial conditions, if quality variability is introduced, the following two propositions hold (proofs are provided in the appendix):

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

Note that Proposition I holds for any learning rule and any quality distribution. The main result of proposition I can be summarised 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 must be monotonously decreasing. Whether in the long-term the market will totally collapse or whether it will get into a stable equilibrium depends on the quality distribution and on the particular learning rules used by buyers.

***Proposition II:***
*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 shown in figures 1, 2 and 3. 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 based only on their own past experience may not seem realistic for those markets in which information can be easily shared between 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).

The market damage caused by quality uncertainty with individual learning is due to the fact that new information about product quality arrives only when there is a new purchase. As lower quality expectations imply lower chances of purchasing a new item, long-sustained low-quality expectations are favoured over long-sustained high-quality expectations (assuming that the learning rule is not biased). The dynamics of quality expectations are asymmetric, because the flow of information, and consequently the probability of updating the expected quality, is conditioned on the particular value of the expected quality.

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

We provide some simulation results from a model of information sharing through randomly generated social networks, where links are created between randomly selected pairs of buyers. Robustness of our results to changes in the network-generating algorithm will be discussed later.

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 (Figure 4 shows representative runs). The general pattern is the same for different quality distributions: uniform, trimmed normal, or exponential (as in the following figures).

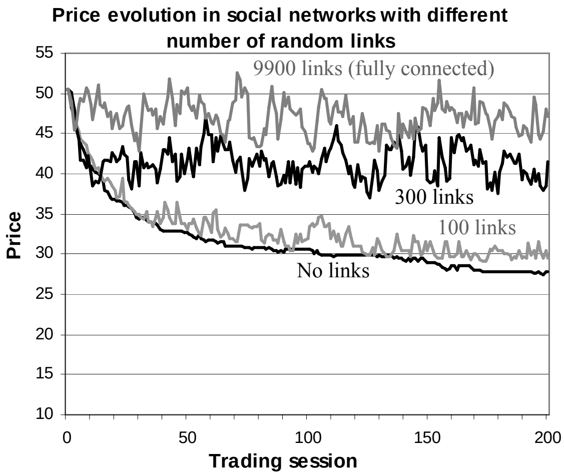


Figure 4. Price evolution in 4 random social networks with 100 buyers, 100 sellers, and different number of random links. Quality distribution $q \sim \exp(1)$, $\lambda_{ind} = 0.25$, $\lambda_{soc} = 0.25$.

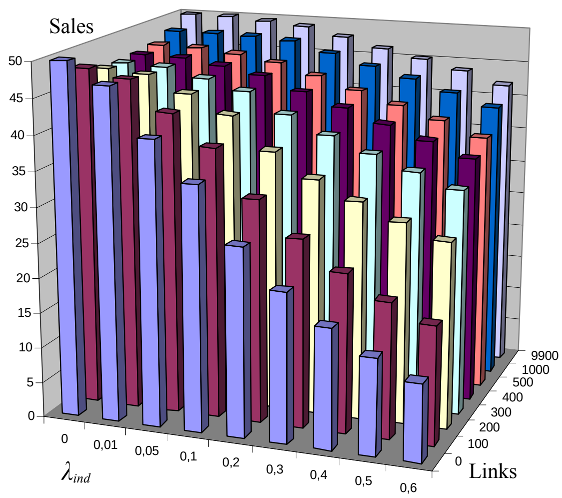


Figure 5. Average (across 1000 random networks in every case) sales at trading session 500, measured in models with different $\lambda_{ind}$ and number of random links, with 100 buyers, 100 sellers, $\lambda_{soc} = 0.4$ and $q \sim \exp(1)$.

Figure 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 $\lambda_{ind}$. There is not much variability across runs (the standard deviation for sales is less than 3.6 units in every case; the standard error for the average values shown in the graph is less than 0.12 in every case). Similar patterns can be observed in average expected qualities, prices and sellers' surplus.

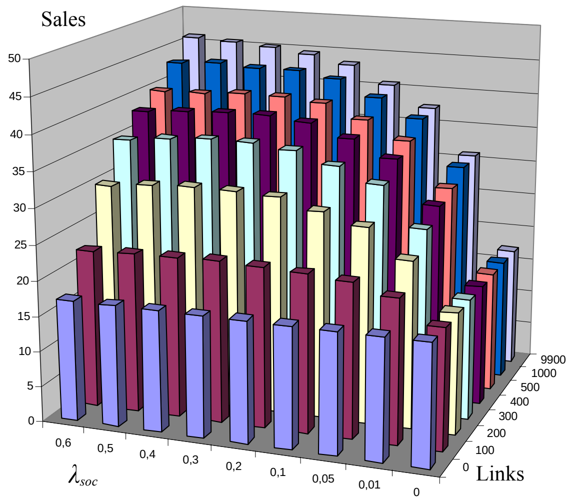


Figure 6. Average (across 1000 random networks in every case) sales at trading session 500, measured in models with different $\lambda_{soc}$ and number of random links, with 100 buyers, 100 sellers, $\lambda_{ind} = 0.4$ and $q \sim \exp(1)$.

Note that for higher $\lambda_{ind}$ there is more variability in quality estimations, causing an effect similar to that of greater quality variability (*i.e.* lower prices and number of sales, more marked quality underestimation, and higher losses in 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 (Figure 5 and Figure 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. Figure 6 shows the (non-linear) effect of $\lambda_{soc}$ keeping the value of $\lambda_{ind}$ constant.

**Robustness to different network structures**

Randomly generated networks can be a good way to test the robustness of a market effect to changes in the network structure (after all, given a certain number of links, any possible network design has a positive probability of being generated by the random procedure we are using). However, different algorithms for network creation will lead to different statistical regularities on the behaviour of the resulting networks.

We tested the validity of our results using some other network-generating algorithms, like the "preferential attachment" rule of Barabási and Albert

(1999) as described by Newman (2003, section VII B). Our general results were robust to changes in the network-generating algorithm, but note, however, that the same network-generating algorithm can give rise to particular networks with very different behaviours. For instance, consider a "star" network-generating algorithm such that one buyer is randomly selected to be the "centre of the star" and a bidirectional link is created between her and each one of the other buyers. The properties of the market in a "star" network critically depend on the behaviour of the buyer in its centre. 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 her market expectations occasionally, between periods of increasing loss of confidence.

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). Reference conditions (no quality variability) for price and sales in this market are close to 50. The evolution of prices (with quality variability) in a "star" social network whose central buyer has a standard reservation price of 63 is shown in Figure 7.

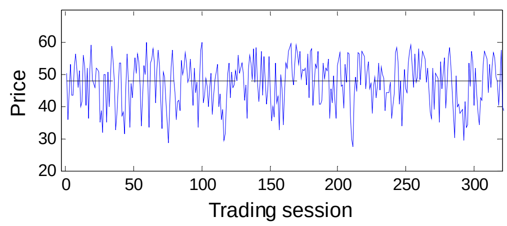


Figure 7. Price evolution in a market model with a "star" social network. The standard reservation value of the central buyer is 63. Conditions: 100 buyers, 100 sellers, $\lambda_{ind} = 0.4$ , $\lambda_{soc} = 0.4$, $q \sim U(0,2)$.

Consider now the same sellers and buyers, also embedded in a "star" social network, but the central buyer has now a standard reservation price of 25. As before, reference conditions (no quality variability) for 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. We can observe the evolution of prices for one of these networks in Figure 8, with periods of loss of confidence going on between shocks caused by purchases of the central buyer. Shocks are usually upwards because new items for the central buyer come with a quality which is usually above the (depressed) average expectations. After a price recovery the central buyer stops buying until prices go down to the level of her reservation price. Thus, it is shown that the same (stochastic) network-generating algorithm can lead to specific networks with dramatically different behaviour.

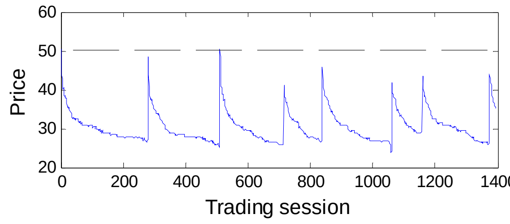


Figure 8. Price evolution in a market model with a "star" social network. The standard reservation value of the central buyer is 25. Conditions: 100 buyers, 100 sellers, $\lambda_{ind} = 0.4$ , $\lambda_{soc} = 0.4$, $q \sim U(0,2)$.

# 4 Discussion and conclusions

The objective of this paper is to analyse the impact of quality variability on markets. This analysis has classically been carried out in the particular framework of asymmetric information and adverse selection. While 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. In this paper we have investigated the effect of quality uncertainty in a more general framework where information is not necessarily asymmetric and buyers estimate product quality using past experiences.

Considering this framework, our model recognises that quality expectations may not be common to every buyer, but instead they may depend on buyers' personal experiences with the product. We then assume one single homogeneous quality distribution for every item. This last assumption is not instrumental to observe the effects of quality variability that we are discussing, but it 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 that we illustrate in this paper is that quality variability combined with the assumption that buyers estimate product quality using their past experience can significantly damage the market, and 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 our 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, and consequently less likely to update their low-quality perception of the product, than those buyers who happen to receive a higher quality item. Thus, low quality expectations tend to persist for longer than high quality ones. New purchases carry 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 we also show 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 we have discussed 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 we are showing with our model 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 analysing a market, the aggregate results of the loss-of-confidence effects we have 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 quality differences between 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 we are describing or because of buyers' risk aversion.

- The case for adverse selection is also weak when there are few agents in the market and repeated 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. The average quality may be easily calculated in markets where individual buyers can check the quality of a large number of items (*e.g.* insurance companies), but otherwise it may not be so (*e.g.* used cars markets).

- In a given market, the importance of personal past experiences (as apposed to aggregate indicators) in people's purchasing behaviour can be empirically tested, either through surveys or through controlled experiments.

- The explanation we put forward in this paper predicts average expected quality to be lower than real average quality, but the model based on asymmetric information assumes that the real average quality is known to every buyer. Thus, detecting such a difference between real and perceived quality would be an indication of the potential presence of the effect we have investigated here.

Finally, note that we have been discussing a loss-of confidence effect due to quality variability at industry level and assuming product homogeneity. A somewhat related situation is that of different firms who provide items with similar average quality but 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 consisting in reimbursing the cost of any defective item *and* replacing it with a new one. The rationale behind this policy is not only reassuring buyers' *a priori* confidence on the product's quality (a cheap good warranty is a clear signal of good quality), but also restoring the buyer's liking for the product if she/he happens to get a defective item, and prevent her/him from switching to another brand.

# Acknowledgements

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

# Appendix:

## Proof of Proposition I

Let [$msp^1$, $msp^2$, …] be the vector of minimum selling prices of the items in the market, sorted out in ascending order, and let $[R^1_{\bullet,n}, R^2_{\bullet,n}, ..., R^{num-buyers}_{\bullet,n}]$ be the vector of *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*. Note that, given the definition of $v_n$, for any number of units *i* with $0 < i <$ *num-buyers*, it holds that:

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

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

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

At session *n* there are (*num-buyers* – $v_n$) buyers who do not purchase any item and whose reservation prices are no greater than $R^{v_n+1}_{\bullet,n}$. As those buyers will not change their reservation price for the next session, we have $R^{v_n+1}_{\bullet,n+1} \le R^{v_n+1}_{\bullet,n}$, from where, using (2), $R^{v_n+1}_{\bullet,n+1} < msp^{v_n+1}$, and using (1), $v_{n+1} < v_n + 1$, which implies the result we wanted to prove:

$$v_{n+1} \le v_n.$$

## Proof of proposition II:

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

$$Hmsp_n = msp^{v_n}$$

Let us divide the set of buyers in a given session *n* into two subgroups:

- Subgroup "Potential purchasers": Those buyers with reservation prices greater than or equal to $msp^{v_n}$. Every purchaser in a session must be in this subgroup.

- Subgroup "Outsiders": 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 there is some trade, 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 (which is what we are trying to prove), the highest minimum selling price remains equal to $msp^{v_n}$. This means that while the number of traded units remains equal to $v_n$ , individuals in the group "outsiders" will not be able to purchase any item, and will therefore stay in that group. On the other hand, individuals in the group "potential purchasers" may move to the group "outsiders", 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 $Hmsp_m = msp^{v_m} = msp^{v_n}$. When the number of individuals in the group "potential purchasers" drops below $v_n$, then the number of traded units will necessarily decrease.

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

## Note on the quality expectations updating rule of section 2.6

Note that the additive model of section 2.6:

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

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

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

$$\lambda_{ind} = \alpha_1 \cdot \alpha_2 \qquad \lambda_{soc} = \alpha_1 \cdot (1-\alpha_2)$$

and it is also equivalent to a model in which the individual (social) observation modifies the expected quality first and then the social (individual) observation modifies the new expectation:

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

$$\lambda_{ind} = \alpha_{ind} \cdot (1-\alpha_{soc}) \qquad \lambda_{soc} = \alpha_{soc}$$

## References


Akerlof GA. The Market for Lemons: Quality Uncertainty and the Market Mechanism. Quarterly Journal of Economics 1970; 84: 488-500.

Barabási AL. Albert R. Emergence of scaling in random networks. Science 1999; 286: 509–512.

Bergemann D. Valimaki J. Learning and strategic pricing. Econometrica 1996; 64 (5): 1125-1149.

Ellison G. Fudenberg D. Word of Mouth Communication and Social Learning. Quarterly Journal of Economics 1995; 110 (1): 93-125.

Hendel I. Lizzeri A. Adverse selection in Durable Goods Markets. American Economic Review 1999; 89 (5): 1097-1115.

Izquierdo SS. Izquierdo LR. Galán JM. Hernández C. Market Failure caused by Quality Uncertainty. In: Mathieu P., Beaufils B. and Brandouy O. (Eds), Artificial Economics - Lecture Notes in Economics and Mathematical Systems 564. Springer-Verlag, Berlin, 2005.

Johnson JP. Myatt DP. Multiproduct quality competition: Fighting brands and product line pruning. American Economic Review 2003; 93 (3): 748-774.

Kirman AP. Vriend NJ. Evolving market structure: An ACE model of price dispersion and loyalty. Journal of Economic Dynamics and Control 2001; 25 (3-4): 459-502.

Macho-Stadler I. Pérez-Castrillo JD. An Introduction to the Economics of Information. Incentives and Contracts. Oxford University Press, 2001 (Second edition)

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

Rose C. Equilibrium and Adverse Selection. The RAND Journal of Economics 1993; 24 (4): 559-569.

Smallwood DE. Conlisk J. Product Quality in Markets Where Consumers Are Imperfectly Informed. Quarterly Journal of Economics 1979; 93 (1): 1-23.

Spence AM. Job Market Signaling. Quarterly Journal of Economics 1973; 83: 355-377.

Stigler GJ. The Economics of Information. Journal of Political Economy 1961; 69: 213-225.

Stiglitz JE. The Contributions of the Economics of Information to Twentieth Century Economics. The Quarterly Journal of Economics 2000; 115 (4): 1441 – 1478.

Vriend N. An Illustration of the Essential Difference Between Individual and Social Learning, and its Consequence for Computational Analyses. Journal of Economic Dynamics and Control 2000; 24: 1-19.

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

Wilson CA. Equilibrium and adverse selection. American Economic Review 1979; 69: 313-317.

Wilson CA. The Nature of Equilibrium in Markets with Adverse Selection. Bell Journal of Economics 1980; 11: 108-130.