Market failure caused by quality uncertainty

Segismundo S. Izquierdo[1], Luis R. Izquierdo[2], José M. Galán[3], Cesáreo Hernández[1]

[1]University of Valladolid (Spain)
[2]The Macaulay Institute (Aberdeen, UK)
[3]University of Burgos (Spain)

Abstract

The classical argument used to explain why markets can fail when there is product quality variability (e.g. the used car market) relies heavily on the presence of asymmetric information –i.e. there must exist some reliable quality indicators that can be observed by sellers, but not by buyers. Using computer simulation, this paper illustrates how such market failures can occur even in the absence of asymmetric information. The mere assumption that buyers estimate the quality of the product they buy using their past experience in previous purchases is enough to observe prices drop, market efficiency losses, and systematic underestimation of actual product quality. This alternative explanation is shown to be valid for a very wide range of learning rules and in various market contexts.

1 Introduction

In this paper we investigate the impact of product quality variability in markets. The discussion will be based on several computer agent-based market simulations.

The main features of the artificial market we are considering are: (a) There is only one type of product, whose quality follows a specific, predetermined, probability distribution, (b) buyers in the market estimate the quality of the product following simple rules based on their own experience, and (c) a buyer's product valuation (or reservation price) is proportional to the quality she expects to get. With these conditions, it is shown in this paper that, when (*symmetric)* quality variability is introduced or increased, market prices fall down and the average quality of the product expected by buyers drops below the real average

Note: For a published version from this working paper see

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

For further work on this issue, including an analysis of the effect of social networks, see the Netlogo model at

http://research.izqui.org
http://www.insisoc.org/research/quality

quality of the product in the market. Consequently, market efficiency[1] is also reduced.

These results emerge from the market interactions among individuals trying to learn the average quality of the product they are repeatedly buying, and are robust to changes in the learning rule and market design. No risk-aversion rule is implemented. Although both the real quality distribution and the buyers' learning rules are unbiased, when several agents get together and interact through a market, an asymmetric-bending effect on market demand emerges, which causes prices to fall down and reduces the efficiency of the market.

Our results suggest an alternative theory to explain the classical and well-documented problem of market failure caused by asymmetric information (Akerlof 1970). The classical explanation assumes a pressing supply of low-quality items that lowers the average quality of the product in the market ("adverse selection"). However, in our model, even letting the average product quality remain constant, a market failure, caused by the combination of quality variability plus individual learning, emerges.

As to the practical implications, our theory could explain the success of warranties and quality variability reduction policies in markets in which asymmetric information does not seem to be a clear issue. However, as is often the case, there can be alternative theories that could also explain the same aggregate effects, and this theory is still to be tested empirically.

2 Perfect competition and quality uncertainty

The classical model of perfect competition has proven to be a useful framework for a large number of real markets. The model establishes that, under certain hypotheses, free trade will produce an equilibrium market price and a traded volume at the crossing point of supply and demand. From an efficiency point of view, this is also the optimal production level and production distribution that a theoretical central planner with all the information should choose.

The hypotheses of perfect competition are quite restrictive: product homogeneity, multiple buyers and sellers, perfect information and profit-maximizing agents. However, for many common market institutions, the model has proven to be robust to deviations in several of these hypotheses. For instance, up to certain limits, the efficiency of double auctions has proven to be robust when the hypotheses of perfect information (Smith 1962), large number of players, and players' cognitive capabilities (Gode and Sunder 1993; Bergstrom 2003; Duffy

[1] Market efficiency is a measure of social welfare. It is the sum of sellers' surplus and buyers' surplus. When a transaction between a buyer and a seller is made, the seller's surplus is the difference between the price of the item (income) and the item's marginal cost; the buyer's surplus is the difference between the maximum price that she would have paid for the item (reservation price, or marginal value) and the price actually paid (cost).

2005) are relaxed (the last two references also discuss some other market institutions).

The hypothesis of product homogeneity is strongly related to the problem of quality distribution. Note that, in general, the actual quality of an item is a random variable, and it can only be observed when the product is used or consumed (think, for instance, of any product with a variable life service, like a light bulb). However, it is often the case that a probability distribution for the quality of the product may be known or estimated a priori. We will assume that a product is homogeneous if any two items of that product have the same quality probability distribution.

Breaches in the hypothesis of product homogeneity combined with asymmetric information have been the subject of intense economic research: “The Market for Lemons: Quality Uncertainty and the Market Mechanism” (Akerlof 1970), is thought to be (by the 2001 Nobel prize commission) the single most important study in the literature on economics of information (The Royal Swedish Academy of Sciences 2001). It provides a fruitful framework for the analysis of many real markets, like those of insurance policies and used cars. However, besides product heterogeneity, this framework relays heavily on the existence of some reliable quality indicators that can be observed by sellers, but not by buyers –i.e. asymmetric information. With less restrictive (more general) conditions, our study provides a complementary explanation for market failure when quality uncertainty is present.

The classical explanation of market failure caused by asymmetric information rests on the phenomenon of “adverse selection”, which can be explained along these lines: there are high-quality and low-quality items, but buyers can not distinguish quality when purchasing, so all items are sold at the same price; for a seller, a low quality-item is more profitable than a high-quality one, so the market is flooded with low-quality items; the reduction of the average quality lowers quality expectations, demand and prices, making high-quality items even unprofitable; this ‘diminishing quality O diminishing price’ vicious cycle can go on to the extent of destroying the market.

In contrast to this explanation, which requires (and mixes) the effects of asymmetric information and product heterogeneity, we do not impose any sort of asymmetry or a breach in the hypothesis of product homogeneity (though we do assume that there is a quality probability distribution). In this setup, our paper discusses the effects of quality variability on the market, assuming that buyers follow simple learning rules to estimate the quality of a product.

3 Design of the experiments

As a guiding line for our discussion we will present the results of several computer simulated markets with individuals who form expectations on the quality of the product using their own past experience.

The aggregate effect can be more easily understood using a simplified model in which a market institution produces a price and a traded quantity at the crossing point of supply and demand (like a Walrasian tatonnement), though our results are also shown to be robust to other market mechanisms.

The main features of our simplified model are:

- Buyers and sellers trade in sessions. In each session, each buyer can buy at most one product ('single–item demand').
- The quality q of every produced product follows a symmetric distribution centred on 1.
- Supply is a linear function and does not vary from one session to the next.
- Demand is formed by summing up the individual reservation prices of buyers. Initial reservation prices are such that the initial demand is linear. Buyers' reservation price then varies according to their quality expectations. The reservation price of buyer i in session n is equal to her initial reservation price multiplied by her current estimated quality ($\hat{q}_{i,n}$) for the product.
- In each session, the market is centrally cleared at the crossing point of supply and demand, and all the buyers who have bought a product update their quality expectations according to their experience with the product just bought. In particular, buyers (indexed in i) use the following updating rule:

$$\hat{q}_{i,n+1} = (1-\lambda)\cdot\hat{q}_{i,n} + \lambda\cdot q \qquad (3.1)$$

with an initial estimate $\hat{q}_{i,0} = \mathrm{E}(q) = 1$. Note that λ (learning rate) measures the responsiveness of buyers' quality estimates to new data. Note also that this is an individual learning rule (Vriend 2000), as each buyer's quality estimates are based only upon her own past experience.

This simple model is enough to illustrate a market failure caused by quality uncertainty. For now, we will not discuss whether the agents' learning rules in our model are realistic or not, as our results will be shown later to be valid for much wider class of learning rules. As a matter of fact, following the Keep-It-Simple principle (Axelrod 1997), we restrict the use of our model to its capacity to illustrate and give insight into the global implications of individual decision rules, and we believe that this is performed best by using simple, tractable and robust models.

4 Results and discussion

Using the simplified model described on the previous section, consider a market with an initial situation ($t = 0$) like the one shown on Figure 1, which corresponds to the following parameterisation: there are 200 buyers, and buyer i ($i = 1, 2, 3, \ldots, 200$) has initial reservation price equal to i; thus the initial demand is such that at price p ($p \leq 200$), the number of products demanded is the integer part of $(201 - p)$. In each session, the number of items offered at price p (supply function) is the integer part of p. The market price is taken to be the average ask-bid price for the last traded unit (crossing point of supply and demand). Therefore reference conditions (i.e. no quality variability) are: price = 100.5, traded volume = 100. These conditions would be indefinitely maintained if there was no product variability, or if the learning rate λ was equal to zero.

We now introduce quality variability and individual quality learning. Surprisingly, in our model with symmetric quality variability, inefficient market dynamics emerge, prices drop below reference conditions, and buyers systematically underestimate the actual quality of the product.

We investigated the robustness of our results using different quality symmetric distributions (uniform, triangular, trimmed normal), obtaining the same patterns for all of them. Robustness to the learning rule will be discussed later.

Figure 1 shows some results corresponding to a uniform quality distribution $q \sim$ U[0, 2] and a learning rate $\lambda = 0.3$, with every buyer's initial quality estimate $\hat{q}_{i,0}$ equal to 1. The degeneration of the demand function can be clearly seen in the first 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 we will show later, with these conditions and given enough time, no trading would eventually take place.

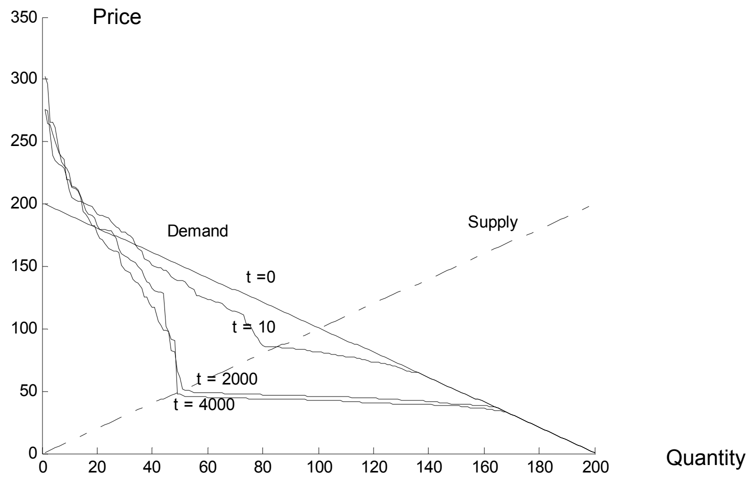

Fig. 1. Effects of quality learning on demand. The quality distribution function for this graph is a uniform distribution U[0, 2]. The initial linear demand (t = 0) is represented, together with the demand functions after 10, 2000, and 4000 trading periods, when changes occur very slowly.

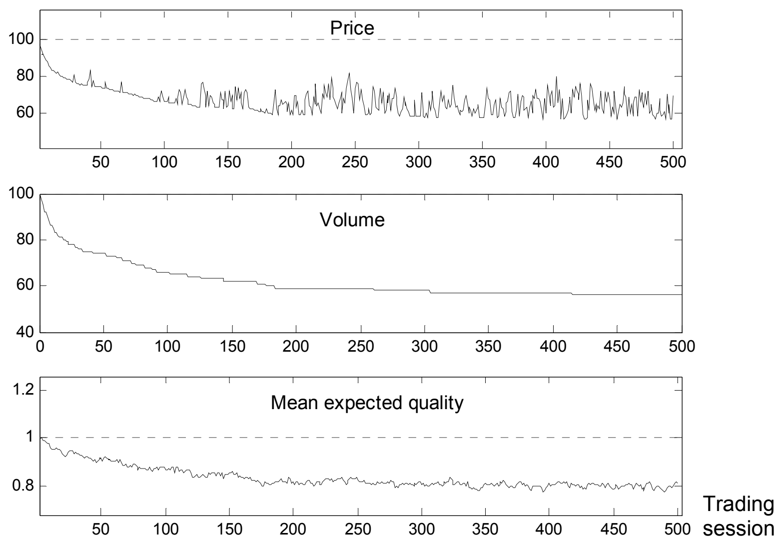

Fig. 2. Effects of quality variability on price level (top), traded volume (middle) and mean expected quality (bottom). The reference situation (without quality variability) is a price level of 100.5, a traded volume of 100 and an expected quality of 1.

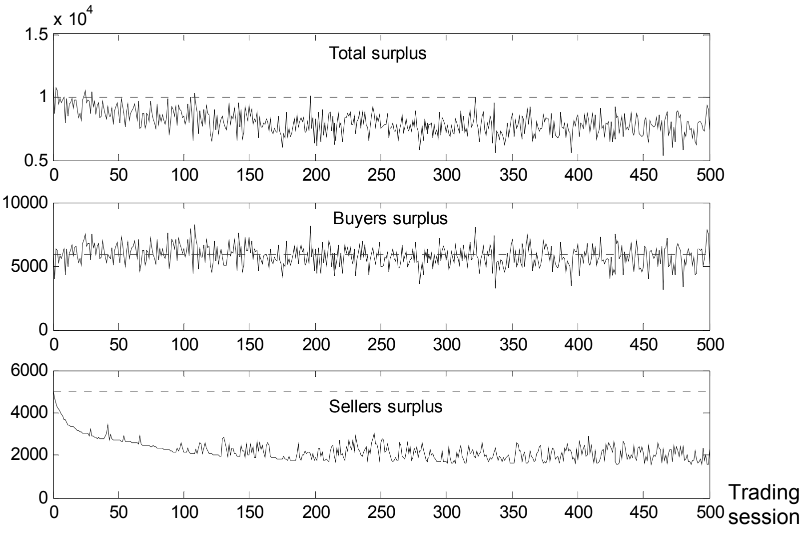

Fig. 3. Effects of quality variability on total surplus (top), buyers surplus (middle) and sellers surplus (bottom). The reference situation (without quality variability) is a total surplus of 10,000: 5,000 for buyers and 5,000 for sellers.

Figure 2 shows the corresponding evolution of prices, traded volumes and average expected qualities. All these variables drop below reference conditions. In Figure 3 we can see the corresponding evolution of surpluses, and the loss of efficiency can be appreciated. Because of the drop in prices, the greatest loss of surplus (benefits) is for sellers.

The general pattern (decreasing prices, decreasing expected quality, monotonously decreasing number of traded units, and loss of efficiency) is consistent throughout simulations, and also for different values of λ and for the different quality distributions considered.

A surprising result of our model is that the average quality of the product in the market is constant, but the average perceived quality is lower than the real one; in fact, most buyers perceive a lower quality than the real one, even though their learning rule is not biased.

The reason for this apparent paradox is that those buyers who at some period(s) get bad products and reduce their quality expectations may stop buying the product, either temporarily (if prices go down enough) or even permanently, remaining forever with low expectations. These low expectations will never increase again if these individuals cannot buy the product once more: their confidence in the product would be undermined forever.

The phenomenon is more clearly seen if we assume that supply is horizontal at a given price level, let us say 50 € (products are sold at 50 €, but not below). If by purchasing a series of “bad” products a buyer’s reservation price can drop below 50 €, she will stop buying the product. In fact, any learning rule that allows quality expectations to occasionally drop below that threshold will have the same effect, finally leading to a market collapse.

More generally, and for any learning rule, if supply is constant and those buyers who do not purchase the product do not change their reservation prices, then the number of traded units must always be monotonous decreasing (this proposition can be proved by noting that, in these conditions, the extramarginal part of the demand, i.e., that part to the right of the crossing point with supply, can never ascend).

Furthermore, in these conditions, let MaC_t be the highest marginal cost of all the units traded at session t, i.e., the marginal cost of the last unit traded. If the combination of learning rule and quality distribution is such that, on any session, there is a positive probability of (after a number of sessions) some reservation price(s) dropping below MaC_t, then the market will eventually collapse. Note that this is the case for our simplified model, because every reserve price may fall below the minimum marginal cost.

5 Robustness to the market mechanism

In this section we check the robustness of our results to changes in the market mechanism by implementing a continuous double auction, as used in stock markets and many other economic institutions (Duffy 2005).

Following Gode and Sunder (1993), we selected zero-intelligence-constrained (ZI-C) agents: buyers who place random bids between 0 and their reservation prices, and sellers who place random asks between their (marginal) costs and a maximum limit (see the model details below). The convergence properties of double auctions with ZI-C agents have been analysed by Cliff and Bruten (1997). They also propose some modifications of the ZI-C agents, like the ZIP agents, that could be argued to be better representatives of human behaviour (Das et al. 2001), but ZI-C agents provide a good starting point to test the effects of a given market measure on the evolution of a double auction, and simple modifications of these agents have been applied to understand phenomena like asset price bubbles and crashes as observed in laboratory market experiments (Duffy and Ünver 2006).

The main features of this continuous double auction model are:

- There are 5 buyers and 5 sellers, sorted out in a queue alternating buyer with seller. They are sequentially prompted to post an offer (ask or bid). Starting with the first agent (buyer or seller), a round is completed when the last (10^{th}) agent is prompted to post an offer.
- Marginal costs for seller j are $MaC_j = 1 + 5\ n_j$, where n_j is the number of units sold by seller j during an auction.
- The quality q of every produced product follows a symmetric distribution centred on 1.
- Reserve prices for buyer i are $R_i = \hat{q}_i\ (200 - 5\ n_i)$, where n_i is the number of units purchased by buyer i during an auction and $\hat{q}_i$ is her current estimated quality (with $\hat{q}_i = 1$ at the beginning of an auction). After every purchase, the quality estimate is updated using Eq. (3.1).
- When prompted for an offer, seller j posts a random ask in the interval [MaC_j, 200], and buyer i posts a random bid in the interval [0, R_i]. The best (highest) bid and the best (lowest) ask are centrally kept, and whenever there is a match (best bid higher than or equal to best ask), a transaction is made between the corresponding buyer and seller. The price of the item is the best bid or the best ask, whichever was first posted.
- After a transaction is made, past offers are cleared and the auction goes on for the next item, continuing the round where it was stopped. An auction ends when there are 100 rounds on end without transactions (when, after a round with transactions, every agent is prompted again for 100 times and no new transaction is made).

A surprising result of double auctions with ZI-C agents (who just make random offers in a profitable range) is the high economic efficiency that can be obtained, very close to the maximum achievable and also very close to the efficiency obtained in experiments with human subjects (Gode and Sunder 1993). Besides, under certain conditions (Cliff and Bruten 1997), the price level tends, as the auction evolves, to the level predicted by perfect competition (crossing point of supply and demand).

In our double auction model, the reference conditions (crossing point of supply and demand with λ =0) are 100 units traded and a price close to 100. Table 1 shows the final price (price of the last unit traded), number of units traded, and efficiency for several auctions before introducing quality variability. Table 2 shows the equivalent results when quality variability (q ~ U[0,2]) and learning (λ = 0.5) were introduced.

Table 1. Results of several auctions, without quality variability

	Auction number						
	1	2	3	4	5	5-105 Average	105-305 Average
Final price	102.8	99.3	102.1	104.1	96.8	100.7	100.4
Units	100	98	98	101	97	99.6	99.7
Efficiency	100%	99.8%	99.8%	100%	99.6%	99.8%	99.8%

Table 2. Results of several auctions, with quality variability q ~ U[0,2] and learning (λ = 0.5)

	Auction number						
	1	2	3	4	5	5-105 Average	105-305 Average
Final price	73.9	87.7	88.7	83.6	68.6	83.3	84.2
Units	78	82	88	88	72	79.8	80.8
Efficiency	77.5%	87.2%	99.4%	92.5%	78.3%	90.5%	91.4%

Table 3. Estimated quality after closing. Results of several auctions with quality variability q ~ U[0,2] and learning (λ = 0.5)

	Auction number					
Buyer number	1	2	3	4	5	1-100 Average
1	0.41	0.82	0.30	0.95	0.46	0.59
2	0.78	0.63	0.59	0.33	0.87	0.58
3	0.71	0.40	0.29	0.60	0.37	0.53
4	1.50	0.75	0.55	0.35	0.52	0.62
5	0.80	1.06	0.43	0.55	0.35	0.58
Average	0.84	0.73	0.43	0.56	0.51	0.58

In short, when quality variability and individual learning were introduced in this double auction environment with ZI-C agents, the general patterns were lower prices, lower perceived quality, less trading, and loss of efficiency. Besides, at the end of most auctions, every buyer underestimates the real quality of the traded product (Table 3).

These results show that the market failure we are describing also happens in a continuous double auction environment: it is robust to changes in the market mechanism.

Intuitively, in this double auction model most buyers finally underestimate the real quality because low expectations have an "attracting" power: the lower the quality expectation, the lower the probability of purchasing a new product and changing the quality expectation. When expectations get low, buying and, consequently, learning, are reduced or even stopped. In fact, it can be graphically checked that, in this model, the demand function tends to get depressed as an auction evolves.

5 Conclusions

We have studied the effect of quality uncertainty in markets using a simple agent-based computer model to gain insight into aggregate market behaviour. We have shown the emergence of a market failure that is not due to buyers' risk aversion, but to some buyers occasionally forming long-lasting low quality expectations, and even leaving the market. The results have proven to be robust to different market institutions, quality distributions and individual learning rules.

Our findings offer a new viewpoint and a new explanation for some well-known aggregate market effects associated to quality variability, as well as to the success of some managerial policies, like the use of warranties or the "zero defects" policy. These effects have traditionally been explained assuming large differences in average quality (product heterogeneity) and asymmetric information, leading to adverse selection. In contrast, our model only assumes learning in an environment of quality uncertainty, which would lead to the same aggregate effects. In order to select between these competing theories, some of the consequences of this model, like the expected difference between average perceived quality and average real quality, are falsifiable and can be tested empirically.

References

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

Axelrod R (1997) The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration. Princeton: Princeton University Press.

Bergstrom T (2003) Vernon Smith's Insomnia and the Dawn of Economics as Experimental Science. Scandinavian Journal of Economics 105: 181-205

Cliff D, Bruten J (1997) Minimal-intelligence agents for bargaining behaviours in market environments. Hewlett–Packard Laboratories Technical Report HPL-97-91

Das R, Hanson JE, Kephart JO, Tesauro G (2001) Agent-Human Interactions in the Continuous Double Auction. Proceedings of the International Joint Conferences on Artificial Intelligence (IJCAI-01)

Duffy J (2005) Agent-Based Models and Human Subject Experiments. To appear in: Handbook of Computational Economics, vol 2. Elsevier, Amsterdam

Duffy J, Ünver MU (2006, forthcoming) Asset Price Bubbles and Crashes With Near Zero-Intelligence Traders. Economic Theory 27: 537-563

Gode DK, Sunder S (1993) Allocative Efficiency of Markets with Zero-Intelligence Traders: Markets as a Partial Substitute for Individual Rationality. Journal of Political Economy 101: 119-137

The Royal Swedish Academy of Sciences (2001) Advanced information on the 2001 Bank of Sweden Prize in Economic Sciences in Memory of Alfred Nobel. Available at http://nobelprize.org/economics/laureates/2001/ecoadv.pdf on May 20th 2005.

Smith VL (1962) An Experimental Study of Competitive Market behavior. Journal of Political Economy 70: 111-137

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