1973

A Model of Consumer Behaviour in a Single Market with Incomplete Information

Antoni Bosch-Domenech

20279

Research Report 7312

A MODEL OF CONSUMER BEHAVIOUR IN A SINGLE MARKET WITH INCOMPLETE INFORMATION

by

Antoni Bosch-Domenech

June 1973.

I. Introduction

The theory of equilibrium competitive analysis is an impressive and beautiful achievement of many generations of economists. But no satisfactory explanation exists of how equilibrium is reached, even though the laws of adjustment were already described by Walras (as tâtonnements), and the underlying principle can be traced back at least to Adam Smith. As is well known, this principle states that prices of commodities increase or decrease depending on whether their excess demand is positive or negative. When this principle is incorporated into a competitive model some odd situations appear. If the economic units are price takers that accept their inability to change prices, who changes them, and why does the pressure of the market compel someone to change them? Those logical difficulties have always been acknowledged. What seems a new development of the last fifteen years[1] is the critical attitude towards the further idealization that was introduced to solve those difficulties: the referee who fixes prices according to some variant of the old principle stated above. There is no doubt that the perfectly-competitive-cum-referee framework has been used to discuss equilibrium economics with gratifying results. But the limitations of this paradigm when applied to stability analysis are too well known. Both the referee and the unawareness of the participants of trading in a market in disequilibrium are theoretical idealizations of doubtful usefulness. They have helped to identify and codify problems in the area of equilibrium economics. But, I believe, any paradigm that incorporates them will fare very poorly when used as the frame of mind for dealing with disequilibrium situations--situations whose empirical relevance offers no doubt (they are everyday's economic reality), and whose theoretical elucidation is basic to the assessment of the merits of equilibrium economics.

Stirred, perhaps, by Hahn's dictum that no unifying principle such as maximization seems available when dealing with disequilibrium situations,[2] a great deal of work has been done in building disequilibrium models in which the participants are maximizers.[3] But, as stated, the formal requirements of perfect competition are internally consistent at a cost of a further idealization and then the paradigm does not provide the adequate intellectual tunnel vision[4] to recognize and solve disequilibrium problems. Therefore, if the maximization postulate had to be kept, something had to go. In particular, the assumption of perfect competition was emptied of two of its basic characteristics. First, individuals were believed aware of the disequilibrium situation, and second, they were postulated as having only imperfect information about the relevant data on which to base their decisions. Unfortunately it turned out that to dig under the Marshallian cross when the cross itself was only partially drawn was not a simple task. And so, we witnessed in the last few years the flourishing of a rather appalling variety of models to explain the behaviour of the market participants and the ultimate approach of the market to some kind of equilibrium, based on different degrees of information and rationality assigned to the individual participants.[5] It seems, therefore, necessary sooner or later (see [11], Conclusions) to start building a systematic taxonomy of cases of rational behaviour based on what each participant knows of each other (what buyers know about sellers, what sellers know about buyers, what buyers know about buyers, and what sellers know about other sellers) and what sorts of equilibria result from each case. Still, it seems to me premature to build such a taxonomy when most of the discussion on the participants behaviour has been conducted with the help of too many <u>ad hoc</u> assumptions introduced to mimic, instead of resulting from, the rational behaviour of the participants.

In particular, I believe that some of the assumptions of the behaviour of rational consumers who do not possess complete information about the market can be derived from more simple postulates. And this is my first task. I will make precise what buyers know about the market prices and about sellers' behaviour and I will postulate a maximizing behaviour.[6] The market in which the transactions take place has a very rudimentary institutional framework. Nonetheless, results are obtained which seem to provide a theoretical explanation of some of the basic peculiarities of markets that the standard models have to disregard. But, as I see it, this paper is especially relevant as the first step towards a more ambitious project, in which a similar treatment of sellers behaviour with different amounts of information will be complemented with the results obtained here, to elucidate the market's evolution and its eventual convergence to some equilibrium position. "Ce ne sont pas les perles qui font le collier, c'est le fil."

II. Outline of the Model of Consumer Behaviour

I consider a single market with an indefinite, very large, but, for the time being, fixed number of stores and consumers. At each period of time every firm fixes a price and each consumer visits one store. The price set by a store can only be modified once at the beginning of each period. To put it in another way, we imagine that it takes one period of time for the consumers to visit one store and this is the same time that the stores take to change prices. Every time the consumer visits one store he observes a price and decides whether he buys at this price or whether he postpones the purchase in the hope of finding a store with a cheaper price. Let me assume that searching comports a cost c (which will be convenient to measure in terms of utility lost) per unit of search that for simplicity I will consider constant. Note that since every store can change the price posted at each period of time a consumer who decides not to buy at one store, in a way loses that opportunity forever since he does not have the guarantee that the store will post the same price in ulterior periods.[7] Finally, I suppose that for every consumer we can define a utility function on the price set U(P), strictly decreasing in price.[8] Therefore the decision of the consumer will be based on a comparison between the net utility that he obtains by buying at the observed price and the expected net utility of searching for another store.

III. Consumer Behaviour with Knowledge of the Price Distribution

Let us suppose not only that the consumer is aware that stores may be charging different prices but he believes he knows the distribution function (d.f.) of prices $\phi(P)$[9,10]. I assume that $\phi(P)$ has a finite second moment and that for the relevant set of prices U(P) is finite. The problem that the utility maximizer consumer confronts is: "When, on the basis of my knowledge and in order to maximize utility, should I decide to buy at the current price instead of looking for another price?" Since the consumer believes he knows the d.f. of the prices he can consider the observed prices $P_1, P_2, \ldots$ as independent,[11] identically distributed random variables from the distribution of prices $\phi(P)$. If he decides to buy after a number n of searches, i.e., after observing $P_n = p_n$, his utility will be $U(P_n) - nc$. Therefore the consumer has to find a rule for stopping his search that maximizes $E[U(P_N) - cN]$, where E is the expectation operator and N is the random number of searches that he undertakes under a particular stopping rule.

Before proceeding we have to make sure that such an optimal rule exists; that is, that the consumer is not going to keep searching forever for a better price. Its existence is guaranteed by Lemma 1 since by the assumptions made above $U(P)$ will have a distribution function $F[U(P)]$ with finite variance. Let us call this maximum expected utility (which will be finite) α^*.

The story unfolds in this way. A consumer decides to buy in the market,[12] and enters one store where he observes a price, P. On the basis of this observation and his belief on the price distribution, he will have to decide either to buy at this price or go to another store. Suppose that he decides not to buy. Then his position is exactly the same as that of before entering the store,[13] except for the fact that he has already incurred a cost c. His maximum expected utility is therefore $\alpha^* - c$ and so, after visiting the first store an optimal procedure is to continue searching if $U(P) - c < \alpha^* - c$ (i.e., $U(P) < \alpha^*$), or to buy if $U(P) - c \geq \alpha^* - c$ (i.e., $U(P) \geq \alpha^*$).

The expected utility of this optimal procedure is $E[\max(U(P),\alpha^*)] - c$. And this is by definition α^*. We have then

$$\alpha^* = E[\max U(P), \alpha^*] - c.$$

<u>Proposition 1</u>

When the common distribution function of each of the observed price P_i is known, has a finite second moment, and $U(P)$ is finite for any relevant price, the optimal rule for the utility maximizer consumer in the market described above is to buy as soon as he finds a price p such that $U(p)$ is at least as large as α^*, where α^* is the unique solution of

$$\int_{\alpha}^{\infty} [U(P) - \alpha]\, dF[U(P)] = c$$

Proof[14]

The conditional expected utility given that the consumer decides to buy at the N^{th} observed price is

$$E(\alpha^*|N) = E(U(P_N)) - Nc ,$$

$$E(\alpha^*) = E[E(U(P_N))] - E(N)c$$

$$E(U(P_N)) = E[U(P_N)|U(P_N) \geq \alpha^*, U(P_{N-1}) < \alpha^*, \ldots, U(P_1) < \alpha^*]..$$

Since the observed prices are independent

$$E(U(P_N)) = E[U(P_N)|(P_N) \geq \alpha^*] .$$

Since the observed prices are identically distributed

$$E(U(P_N)|U(P_N) \geq \alpha^*] = E[U(P)|U(P) \geq \alpha^*]$$

therefore,

$$\alpha^* = E[U(P)|U(P) \geq \alpha^*] - E(N)c .$$

$$E[U(P)|U(P) \geq \alpha^*] = \frac{\int_{\alpha^*}^{\infty} U(P)\, dF[U(P)]}{\int_{\alpha^*}^{\infty} dF[U(P)]}$$

E(N) is the expectation of a random variable that expresses the number of searches required to observe the buying price; that is, the price P at which $U(P) \geq \alpha^*$. This random variable has a geometric distribution with parameter $\Pi = \int_{\alpha^*}^{\infty} dF[U(P)]$, that is, the probability of U(P) being at least as large as α^*. Therefore $E(N) = \Pi^{-1}$.

Finally, then,

$$\alpha^* = \frac{\int_{\alpha^*}^{\infty} U(P)\, dF[U(P)] - c}{\int_{\alpha^*}^{\infty} dF[U(P)]}$$

or $$c = \int_{\alpha^*}^{\infty} [U(P) - \alpha^*]\, dF(U(P)) \qquad [1]$$

Let us call $\int_{\alpha*}^{\infty} [U(P) - \alpha*]\, dF[U(P)] \equiv T_{F[U]}(\alpha*)$

Lemma 3 guarantees, in particular, that $T_{F[U]}(\alpha*)$ is a strictly decreasing function of $\alpha*$ and, therefore, $\alpha*$ is unique, completing the proof of the proposition.

The left-hand term of [1] is the marginal cost of searching for another store and the right-hand side is the expected marginal utility of searching for another store. $\alpha*$ is the number that equates marginal cost and expected marginal utility. It is, in addition, the expected utility of following the optimal stopping rule. It is clear then that if at $\overline{P}$, where $U(\overline{P}) = \alpha*$, $D(\overline{P})$ is zero, the consumer will not enter the market; but if $D(\overline{P}) > 0$, the consumer will decide to buy and will not stop searching until he finds a store that charges a price $P \leq \overline{P}$. Since $T(\alpha*)$ is strictly decreasing on $\alpha*$, and $U(P)$ is strictly decreasing on P, as c increases $\overline{P}$ increases also and there will exist a $\bar{c}$ such that for search costs at least as great as $\bar{c}$ the consumer will not enter the market. This is therefore a natural way of determining the cutoff price.[15] In particular, if the consumer believes that prices are distributed normally with mean μ and variance σ^2 and U(P) is a linear function, by Lemma 4,

$$\alpha* = U(\mu) + \bar{\sigma}\, \Psi^{-1}\left(\frac{c}{\bar{\sigma}}\right) \qquad [2]$$

where $\bar{\sigma}^2$ is the variance of F[U(P)] and $\Psi(S) = \int_S^{\infty} (x-S)\, \varphi(x)dx$ where $\varphi(x)$ is the p.d.f. of a standard normal distribution.

Although it is not possible to have a complete description of the market without matching the above discussion with a model of seller behaviour, a few remarks can still be made. In particular, it is clear from [2] that the larger the expected value of the price distribution the smaller α^*, and that the smaller $\bar{\sigma}^2$, the smaller $\Psi^{-1}\left(\frac{c}{\bar{\sigma}}\right)$ and therefore α^*. So, in a natural way, the previous discussion makes room for some sort of non-pricing competition among sellers. Namely, one would expect from the above result that those stores that (due to their location, name, etc.) are likely to be visited first, will advertise with the purpose of convincing the potential customers that prices in average are high and that "steals" are things of the past. Conversely, the other stores will try to convince customers of the opposite. Those results are congenial to commonsense and probably do not teach us anything new about the behaviour of stores. But it is always comforting to find an explanation even (or should I say especially?) for the obvious.

IV. Consumer Behaviour with Imperfect Knowledge of the Price Distribution

In the previous discussion I supposed that the consumer was aware (or believed to be aware) of the price distribution. This is, no doubt, a very strong assumption--not only because it implies the possession of an amount of information that no consumer would be believed to have, but because it was taken to mean that once the consumer accepts a distribution of prices he sticks to it no matter what are the values of the prices observed in successive searches. Instead, we would prefer to imagine a consumer with some initial idea about the price distribution that would be modified as he kept visiting different stores and being informed about the prices quoted.

Specifically, I assume that the consumer believes (and sticks to this belief) that the price distribution is normal[16] with variance $V = 1$.[17] On the other hand, he does not believe to know the mean M of the price distribution but he enters the market with some idea about it. He believes that M is normally distributed with mean μ_o and precision τ_o.[18] For the rest, the market is as described in the previous section. With all this information the consumer enters the market with the intention of buying at the price, and after the number of searches, that maximize his expected utility $E[U(P_N) - cN]$. Therefore, at each period he will have to decide whether he looks for more information about the prices quoted in the market or whether he stops searching and buys at the last price observed. This decision will be based on the price just observed, P (and its corresponding utility), on the mean μ and the precision τ of the current posterior distribution of M,[19] and finally, on the constant cost c of search. Once again we can safely assume that an optimal stopping rule exists since the conditions of Lemma 2 are satisfied. Let me call the expected utility (exclusive of the search cost) from following the optimal procedure when P has just been observed $U^*(P, \mu, \tau)$.

Suppose that the consumer believes that $M \sim N(\mu_o, \tau_o)$. If he now decides to go into another store and observes the price $P(\mu_o, \tau_o) = p(\mu_o, \tau_o)$, by Lemma 5, the posterior distribution of M will be $M \sim N(\mu_1, \tau_1)$, where

$$\mu_1 = \frac{\tau_o \mu_o + p(\mu_o, \tau_o)}{t + 1} \quad \text{and} \quad \tau_1 = \tau_o + 1 \qquad [3]$$

In general, if he decides to go to i other stores, the posterior distribution of M after the last store has been visited will be $M \sim N(\mu_i, \tau_i)$ where

$$\mu_i = \frac{\tau_o \mu_o + \sum_{j=o}^{i-1} p(\mu_j, \tau_j)}{\tau_o + i} \quad \text{and} \quad \tau_i = \tau_o + i \quad \text{and} \quad p(\mu_j, \tau_j) \text{ is}$$

price observed after j-1 more visits to stores. The marginal distribution of the next observed price will be, by Lemma 6 (given that $(P|M) \sim N(M, 1)$ and

$M \sim N(\mu_o, \tau_o)$),

$$P(\mu_o, \tau_o) \sim N(\mu_o, \frac{\tau_o}{\tau_o+1}).$$

And, in general, the marginal distribution of the price observed after visiting i stores will be

$$P(\mu_{i-1}, \tau_{i-1}) \sim N(\mu_{i-1}, \frac{\tau_o+i-1}{\tau_o+i}).$$

Suppose, then, that after visiting i stores and reading the price posted in the i^{th} store, P, the consumer believes that $M \sim N(\mu_i, \tau_i)$. With this in mind, he may decide to buy or, alternatively, to look for another price $p(\mu_i, \tau_i)$. He will buy if $U(P) \geq E[U^*(P(\mu_i, \tau_i), \mu_{i+1}, \tau_{i+1})] - c$ and will continue searching otherwise. Therefore

$$U^*(P, \mu_i, \tau_i) = \max\{U(P), E[U^*(P(\mu_i, \tau_i), \mu_{i+1}, \tau_{i+1}] - c\} \quad [4]$$

Let us call $E[U^*(p(\mu_i, \tau_i), \mu_{i+1}, \tau_{i+1}] - c \equiv \alpha(\mu_i, \tau_i)$.

We can state the following

<u>Proposition 2</u>

When the common distribution function of each of the observed prices is believed to be normal with a known finite variance, and with mean M believed to be distributed normally with mean μ_i and precision τ_i, and the last observed price is P, the optimal rule for the utility maximizer consumer in the market described above is to buy at the price P if U(P) is at least as large as $\alpha(\mu_i, \tau_i)$.

Now, by [4],

$$\alpha(\mu_i, \tau_i) = E[\max\{U(P(\mu_i, \tau_i)), \alpha(\mu_{i+1}, \tau_{i+1})\}] - c$$

therefore,[20]

$$\alpha(\mu_i,\tau_i) = \alpha(\mu_{i+1},\tau_{i+1}) + \int_{\alpha(\mu_{i+1},\tau_{i+1})}^{\infty} [U(P(\mu_i,\tau_i)) - \alpha(\mu_{i+1},\tau_{i+1})]\, dF[U(P(\mu_i,\tau_i))] - c$$

$$\equiv \alpha(\mu_{i+1},\tau_{i+1}) + T_{F[U(P(\mu_i,\tau_i))]}(\alpha(\mu_{i+1},t_{i+1})) - c$$

Since $U(P)$ is known and the d.f. of $P(\mu_i,\tau_i)$ $(i=o,1,\ldots)$ is known for any μ_o,τ_o, the d.f. of $U[P(\mu_i,\tau_i)]$ is known. Since, moreover, $T(\alpha(\mu_{i+1},\tau_{i+1}))$ (where the subscript has been dropped for notational simplicity) is strictly decreasing on $\alpha(\mu_{i+1},\tau_{i+1})$ by Lemma 3, $\alpha(\mu_{i+1},\tau_{i+1})$ is uniquely defined from the knowledge of $\alpha(\mu_i,\tau_i)$ and, in general, any $\alpha(\mu_i,\tau_i)$ $(i=1,2,\ldots)$ can be uniquely defined from the knowledge of $\alpha(\mu_o,\tau_o)$. In particular, if the utility function is <u>linear</u> on prices,

$U[P(\mu_i,\tau_i)] \sim N(m_i,\pi_i)$ where $m_i = U(\mu_i)$ and $\pi_i = b\tau_i$, $b > 0$.

Therefore, $T(\alpha(\mu_{i+1},\tau_{i+1})) = \pi_i^{-\frac{1}{2}} \Psi[\pi_i^{\frac{1}{2}}(\alpha(\mu_{i+1},\tau_{i+1}) - m_i)]$

and $$\alpha(\mu_i,\tau_i) = \alpha(\mu_{i+1},\tau_{i+1}) + \pi_i^{-\frac{1}{2}} \Psi[\pi_i^{\frac{1}{2}}(\alpha(\mu_{i+1},\tau_{i+1}) - m_i)] - c.$$

Using Lemma 8,

$$\alpha(\mu_i,\tau_i) = \pi_i^{\frac{1}{2}} \Psi[\pi_i^{\frac{1}{2}}(m_i - \alpha(\mu_{i+1},\tau_{i+1}))] + m_i - c .$$

Little can be said about $\alpha(\mu_i,\tau_i)$. The a priori mean m_i enters the right hand side of the above equality as an additive term, as an argument of the strictly decreasing function Ψ and in $\alpha(\mu_{i+1},\tau_{i+1})$. But at least it can be verified that the entire sequence of $\alpha(\mu_i,\tau_i)$, $i=0,1,\ldots$, is determined recursively from any of its elements. This is so since Ψ can be inverted to yield:

$$\alpha(\mu_{i+1},\tau_{i+1}) = m_i - \pi_i^{-\frac{1}{2}} \Psi^{-1} [\pi_i^{\frac{1}{2}} (\alpha(\mu_i,\tau_i) - m_i + c)].$$

Yet, since the utility and price spaces are homeomorphic, the decision problem that the consumer has to solve can be discussed directly in terms of utilities. If $Z = U(P)$, Z is a random variable distributed normally with unknown mean $\Lambda = U(M)$ an precision b, a positive constant. The prior distribution of Λ is normal with mean $m_o = U(\mu_o)$ and precision $\pi_o = b\tau_o$. We will designate by $\mathcal{E}(z,m,\pi)$ the expected utility (exclusive of the search cost) from following the optimal procedure when $z = U(p)$ has just been observed and the consumer believes, therefore, that m_i and π_i are the mean and the precision, respectively, of the current posterior distribution of Λ. Of course, $\mathcal{E}(z,m_i,\pi_i) = U^*(p_1,\mu_i,z_i)$. As before

$$\mathcal{E}(z,m_i,\pi_i) = \max \{z, E[\mathcal{E}(Z(m_i,\pi_i), m_{i+1},\pi_{i+1})] - c\} .$$

Now suppose that the present and future observed utilities, as well as the mean of Λ, are reduced by k. The expected optimal utility after observing $z - k$ and believing that $\Lambda \sim N(m_i - k, \pi_i)$ must equal the expected optimal utility with the original data minus k, or

$$\mathcal{E}(z - k, m_i - k, \Pi_i) = \mathcal{E}(z, m_i, \Pi_i) - k . \quad [5]$$

This being so it can be stated that

$$\mathcal{E}(\hat{z},o,\pi_i) = \max \{\hat{z},E[\mathcal{E}(Y(\pi_i), o, \pi_{i+1})] - c\} \quad , \quad \hat{z} = z - m_i$$

where $Y(\pi_i) \sim N(o, \pi_i+b/\pi_i)$[21].

Calling $E(\mathcal{E}(Y(\pi_i), o, \pi_{i+1})) - c = \overline{\alpha} (\pi_i)$,

we can establish that the optimal rule for the consumer that has just observed price $p = U^{-1}(z)$ and believes, therefore, that $\Lambda \sim N(m_i,\Pi_i)$ is to buy at this price if $z \geq \bar{\alpha}(\Pi_i) + m_i$ and to look for another store otherwise.

In this way, (m_i,Π_i) can be expressed as the sum of m_i plus another term, $\bar{\alpha}(\Pi_i)$, independent of m_i. It is clear, therefore, that the higher the expectation of the price mean, the more likely it is that the number of stores visited will decrease, and the lower will be the expected utility. This result will hardly surprise anybody. Less intuitive is the conclusion that the higher the precision of the distribution of the price mean, the more likely it is that the number of stores visited will increase and the higher will be the expected utility. This is so since $\bar{\alpha}(\Pi_i)$ is a continuous, strictly increasing function of Π_i, (see [4], p.339). These results confirm, therefore, that as higher and higher prices are observed, caeteris paribus, the cutoff price, $\bar{p} = U^{-1}(\bar{z})$, $\bar{z} = \bar{\alpha}(\Pi_i) + m_i$, will rise[22].

To obtain further results it is indispensable at this stage to have a close look at $\bar{\alpha}(\Pi_i)$, which is the algebraic sum of two terms. The first one, $E(\mathcal{E}(Y(\Pi_i), 0, \Pi_{i+1}))$, can be interpreted in the following way. Suppose that a consumer searching for prices can "observe" a sequence of utility random variables $Y(\Pi_1), Y(\Pi_2), \ldots$, at a cost c. These random variables are independent but not identically distributed. The distribution of $Y(\Pi_i)$, $i = 1,2,\ldots$, is $Y(\Pi_i) \sim N(0, \Pi_0 + i / \Pi_0 + i - 1)$, for a given, initial Π_0. Then $E(\mathcal{E}(Y(\Pi_i), 0, \Pi_{i+1}))$ is the optimal expected utility when $Y(\Pi_i)$ has been observed. The second term is the search cost.

Therefore the optimal procedure is to buy at the observed price if its utility minus the utility of the a priori expected price mean is greater than the optimal expected utility in the auxiliary search problem described above, minus the search cost. This means that observed prices whose utility is greater than the utility of the a priori expected price mean minus the search cost can be rejected. Or, more surprisingly, that the consumer might prefer to keep visiting stores (at the given cost) than to buy at an observed price whose utility is greater than the utility of the a priori expected price mean. This is not, though, an unnatural result. If the search cost is not very high and the consumer believes that the price distribution is not too concentrated around the mean, since he is not obliged to buy when he enters a store, he might be willing to pay the cost of trying his luck in the search for a better bargain.

A corollary of this result is the implication that we cannot guarantee ever increasing cutoff prices of the consumers staying in the market, since it is conceivable that a consumer observing a price below the a priori expected mean will consider that his chances of finding an even lower price are enhanced[23], staying therefore in the market with a lower cutoff price. But since $\bar{\alpha}(\Pi_i)$ is a continuous, strictly increasing function of Π_i and it can be shown[24] that

$$\lim_{\Pi_i \to \infty} \bar{\alpha}(\Pi_i) = \psi^{-1}(c),$$

and, in addition, $\psi(0) = (2\pi)^{-1/2}$ and therefore $\psi^{-1}((2\pi)^{-1/2}) = 0$,

it follows that

$$c \geq (2\pi)^{-1/2} \quad \leftrightarrow \quad \psi^{-1}(c) = \lim_{\Pi_i \to \infty} \bar{\alpha}(\Pi_i) \leq 0,$$

i.e., $\bar{\alpha}(\Pi_i)$ will be negative for any Π_i if and only if $c > (2\pi)^{-1/2}$.[25]

In other words, given a utility measurement unit, for costs of search lower than $(2\pi)^{-1/2}$, there might exist levels of Π_i high enough such that the cutoff price decreases from one period to another (provided, of course, that the observed price is lower than the a priori expected mean).

Therefore, to guarantee ever increasing cutoff prices of the consumers remaining in the market we will have to make the assumption that $c > (2\pi)^{-1/2}$. If we do not pay too much attention to the specific value of this lower bound (which is simply the result of the assumption of the price and mean distributions being normal) this is not as awkward an assumption as it may seem. It simply states that if the search cost is too low, the consumer might be willing to try his luck with another store even if the price observed is lower than the average price expected. Actually, in the limit, with no search cost, the search would never end.

Finally, let me observe that the distinction that Diamond [5] makes between tourists and residents can be incorporated in this model as one of degree, characterized by the a priori precision of the price distribution mean. Note that it is conceivable that a tourist may have a fair idea of the dispersion of prices in the country that he is visiting, while his knowledge of what price to expect for a certain good might be very incomplete. More likely he only has an approximate idea of the average price in his own country and a scaled version of it might be taken as a priori mean of M. The a priori precision, at any rate, will probably be very low. A resident on the other hand might be considered as having a more precise knowledge of the distribution of M and, when the precision of his knowledge becomes very high, the behaviour of the consumer described in section III may be an adequate approximation of his behaviour.

LEMMATA

The following lemmas are given without proof. They are either proved on the references, well known results, or have proofs which are immediate.

Lemma 1.

Let $X_1, X_2, \ldots$ be a sample of independent, identically distributed random variables from a distribution for which it is known that the d.f. is F. Let c be a fixed cost per observation and for n = 1,2,..., let $Y_n = X_n - nc$. If $E(X_i^2) < \infty$ for i=1,2,..., then there exists a stopping rule which maximizes $E(Y_N)$. (See [4], p. 353).

Lemma 2.

Let $X_1, X_2, \ldots$ be a sample of independent, identically distributed random variables from a distribution which involves a parameter M whose value is unknown, and suppose that M has a specified prior distribution, and for n = 1,2,..., let $Y_n = X_n - nc$. If $E(X_i^2) = E[E(X_i^2|M)] < \infty$ for i=1,2,..., then there exists a stopping rule that maximizes $E[Y_N]$. (See [4], p. 353).

Lemma 3.

Let F be a distribution function on the real line for which a mean μ exists. Then $\int_s^\infty (x-s)\, dF(x)\ -\infty < s, < \infty$ is a non-negative strictly decreasing function of s.

Lemma 4.

Let F be the d.f. of a normal distribution with mean μ and precision τ. Then $\int_s^\infty (x-s)\, dF(x) = \tau^{\frac{1}{2}} \int_{\tau^{\frac{1}{2}}(s-\mu)}^\infty [z - \tau^{\frac{1}{2}}(s-\mu)]\, \varphi(z) dz$ where z is a standard normal random variable and $\varphi(z)$ is its p.d.f.

Lemma 5.

Let $X_1,\ldots,X_n$ be a random sample from a normal distribution with an unknown value of the mean M and a specified value of the precision $r(r > 0)$. Suppose that the prior distribution of M is a normal distribution with mean μ and precision τ such that $-\infty < \mu < \infty$ and $\tau > 0$. Then the posterior distribution of M when $X_i = x_i$ $(i=1,\ldots,n)$ is a normal distribution with mean μ' and precision $\tau + nr$, where

$$\mu' = \frac{\tau\mu + nr\bar{x}}{\tau + nr}$$

(See [4], p. 167).

Lemma 6.

Let X and Y be two random variables. Suppose that the d.f. of $(X|Y = y)$ for $-\infty < y < \infty$ is normal with mean y and precision τ_1 and that the marginal distribution of Y is normal with mean μ and precision τ_2. Then the marginal distribution of X is normal with mean μ and precision $\frac{\tau_1\tau_2}{\tau_1+\tau_2}$.

Lemma 7.

Let X be a random variable with a d.f. F(x) for which the mean exists. Then, for any $s \in (-\infty, \infty)$,

$$E[\max\{X,s\}] = s + \int_s^\infty (x-s)\, dF(x) .$$

(See proof of Theorem 1.)

Lemma 8.

$\int_{-s}^\infty (x+s)\, \varphi(x)\, dx = s + \int_s^\infty (x-s)\, \varphi(x)\, dx$ for any $-\infty < s < \infty$, where $\varphi(\cdot)$ is the p.d.f. of a standard normal distribution.

Footnotes

[1]See [1] p. 43, [7] p. 179 or [2] p. 49.

[2]See [6] p. 1

[3]For some work in disequilibrium models with maximizing agents see [11].

[4]To use an expression of B. J. Loasby in [8].

[5]For a survey of this field see [11].

[6]This follows the traditional methodology of economic theory which basically views the concept of rational behaviour as the rational choice of an isolated individual which maximizes a well defined function.

[7]Actually, if there is a sufficiently large number of stores, the same results would be obtained if it was assumed that the stores were not allowed to modify the prices once posted. In that case if the consumer decided to buy after visiting n stores his utility would be $\max\{U(p_1),\ldots,U(p_n)\} - nc$.
But since the observed prices are independent and identically distributed the consumer does not get additional information about the price distribution that would induce him to buy at a price observed at a previous stage and that was, then, considered too high. About this see [4] p. 335.

[8]I suppose that there is a unique relationship between the price and the quantity that the consumer decides to buy, D(p), independent of the number of searches.

[9]If prices can be changed after each period it is likely that the true distribution of prices will change with time. But the results obtained depend on the consumer's belief of the price distribution which, in this section, is unaffected by the price observed.

[10]$\phi(p)$ expresses the probability of a price, not of a price in a particular store.

[11]Strictly speaking this might be a strong assumption. Still with a large number of indifferentiated stores it seems plausible.

[12]As in [5] we might think of a market of a durable good once the decision of buying is taken the amounts D(p) bought might be considered as representing quality and size differences. About the decision to buy see the end of this section.

[13]Once the consumer has decided not to buy at the price charged by some store and to keep searching for a better deal, the information about prices gathered so far becomes irrelevant since prices may change from period to period. But see note 6.

[14] A similar result has been proved for instance in [9]. I follow [3] for the proof of the text.

[15] This is the largest price at which the consumer will decide to buy.

[16] This seems in general a reasonable assumption if the market is composed of large numbers of firms. What may not be that reasonable is that no matter the observed prices, the consumer sticks to his belief.

[17] To take $V = 1$ is a simplification that does not bring any loss of generality. In addition it has the advantage of converting the variance of the posterior distribution of M in a simple index of the number of stores visited.

[18] The precision is defined as the inverse of the variance. It is used, following [4], since it seems to simplify somewhat the notation.

[19] That is, the mean and the precision that the consumer assigns to the distribution of M after taking into account the value of all the previously observed prices.

[20] See proof of Proposition 1 or Lemma 7.

[21] De Groot [4] ,p. 337-8 gives the following proof.From [3] (with the appropiate notational change,since now we are dealing directly with utilities) and [5],

$$(\hat{z},0,\Pi_i) = \max\{\hat{z}\ ,\ E(\mathcal{E}\ (\hat{Z}(\Pi_i),\ \hat{Z}(\Pi_i)/\Pi_i+1\ ,\ \Pi_i+1)) - c\} =$$

$$= \max\{\hat{z}\ ,\ E(\mathcal{E}\ (\Pi_i Z(\Pi_i)/\Pi_i+1\ ,\ 0,\ \Pi_i+1\ )) - c\}\ .$$

And since $\hat{Z}(\Pi_i) \sim N(0\ ,\Pi_i/\Pi_i+1)$, calling $Y(\Pi_i) = \Pi_i\ Z(\Pi_i)/\Pi_i+1$, the equality follows.

[22] One would be tempted to claim that the proposed model yields the sensible result that if a consumer observes an inflationary trend,the desirability of buying now will increase.Unfortunately,the assumption of independence of prices does not allow for this interpretation.

[23] Diamond [6] points out the possibility of such behaviour,although he observes that it is more likely in the early stages of the search.Stangely enough,our model yields the opposite result,namely,that this behaviour is more likely as more and more stores are visited.This is due to the fact that as the number of stores visited increases,the precision of the distribution of M increases.

[24]See De Groot [4],p. 339.

[25]Note that $(2\pi)^{-1/2}$ are units of utility.This is so since

$$\lim_{\Pi_i \to \infty} \bar{\alpha}(\Pi_i) = \lim_{\Pi_i \to \infty} (\Pi_i / 1+\Pi_i)^{1/2} \Psi^{-1}(c (1+\Pi_i/\Pi_i)^{1/2}).$$

This limit will be zero whenever

$$\lim_{\Pi_i \to \infty} \Psi^{-1}(c(1+\Pi_i/\Pi_i)^{1/2}) = 0 ,$$

i.e.,whenever $\lim_{\Pi_i \to \infty} c(1+\Pi_i/\Pi_i)^{1/2} = (2\pi)^{-1/2}$,which is to say whenever

$$\lim_{\Pi_i \to \infty} (2\pi)^{-1/2} (\Pi_i/1+\Pi_i)^{1/2} = (2\pi)^{-1/2} = c ,$$

where both sides of the equalities have the same units,since both c and Π_i are measured in units of utility.

[26]Diamond [6] observes that the presence of tourists may lower the price for residents.Evidently,whether this is true will depend on the sort of store behaviour postulated.But on the light of the above results one might conjecture that,if the stores behave "rationally",the larger the number of consumers with imprecise a priori beliefs about the price mean (we have called them tourists),the larger will be the market prices.

References

[1] Arrow, K. J., "Toward a Theory of Price Adjustment," in The Allocation of Economic Resources, by M. Abramovitz et al., Stanford University Press, Stanford, 1959.

[2] Chipman, J. S., "The Nature and Meaning of Equilibrium in Economic Theory," in Functionalism on the Social Sciences, Monograph No. 5 of the American Academy of Political and Social Science, Philadelphia, February 1965.

[3] Chow, Y. S. and Robbins H., "A Martingale System Theorem and Application, Proc. 4th Berkeley Symp. Math. Stat. and Prob., 1961.

[4] De Groot, M. H., Optimal Statistical Decisions, New York, McGraw-Hill, 1970.

[5] Diamond, P. A., "A Model of Price Adjustment," Journal of Economic Theory, 3, pp. 156-168, 1971.

[6] Hahn, F. H., "Some Adjustment Problems," Econometrica, 38-1, January 1970.

[7] Koopmans, T. C., Three Essays on the State of Economic Science, New York, McGraw-Hill, 1957.

[8] Loasby, B. J., "Hypothesis and Paradigm in the Theory of the Firm," Economic Journal, December 1971.

[9] MacQueen, J. and Miller, R. G., Jr., "Optimal Persistence Policies," Operations Research, 8, 1960.

[10] Rothschild, M., "Models of Market Organization with Imperfect Information: A Survey," unpublished paper.