



The Nature of Equilibrium in Markets with Adverse Selection

Charles Wilson

The Bell Journal of Economics, Vol. 11, No. 1 (Spring, 1980), 108-130.

Stable URL:

<http://links.jstor.org/sici?sici=0361-915X%28198021%2911%3A1%3C108%3ATNOEIM%3E2.0.CO%3B2-2>

The Bell Journal of Economics is currently published by The RAND Corporation.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/rand.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

The nature of equilibrium in markets with adverse selection

Charles Wilson*

In the presence of adverse selection, how does the nature of the market equilibrium depend on the convention used to set the prices? Using a variant of Akerlof's model of the used car market, we examine the equilibrium of the model under three distinct conventions: (1) an auctioneer sets the price; (2) buyers set the price; (3) sellers set the price. Only in the case of the auctioneer is the equilibrium necessarily characterized by a single price which equates supply and demand. When either buyers or sellers set the price, a distribution of prices may emerge with excess supply at some or all of the prices. The analysis suggests that the allocation of goods in markets where adverse selection is a serious problem may be sensitive to the convention by which prices are set.

1. Introduction

■ Consider the following version of Akerlof's (1970) familiar paradigm of a market with adverse selection. A set of owners wish to sell their used cars to a set of potential buyers. Sellers differ in the quality of car they own. Buyers differ in the value they attach to cars of the same quality. Each seller knows the quality of his own car; each buyer, however, can observe only the average quality of the cars sold at each price. What is the nature of the market equilibrium? The conventional answer to this question presumes that if trade takes place at all, it must take place at a single price. A distribution of prices is possible only if there is some other observable characteristic correlated with quality which can serve as a signal. Furthermore, the mechanism by which prices are formed does not typically play a role in the analysis. Equilibrium is simply defined as a single price which equates supply and demand.

In this paper I argue that the conventional analysis is misleading in two respects. First, in the absence of search costs, the nature of the equilibrium is very sensitive to the price-setting convention adopted by the agents in a market. Unlike the corresponding market with perfect information, it matters whether buyers, sellers, or an auctioneer sets the price. Second, equilibrium need not be characterized by a single price. Depending on how prices are formed,

* University of Wisconsin, Madison.

This research was supported by the National Science Foundation under grants SOC 77-08568 and 78-08980. Space did not permit a complete demonstration of all of the results cited in the text. A more thorough and technical presentation of some of these results appears in an appendix accompanying the unpublished version of this paper and can be made available upon request.

there may be a distribution of prices in equilibrium with excess supply at some or all of these prices. These results suggest that when analyzing markets where the problem of adverse selection is potentially significant, it is important that we understand the precise mechanism through which the prices are set. This may affect not only the outcome we expect to observe, but the welfare implications as well.

To focus on the performance of a market under different price-setting conventions, I shall confine the analysis to the specific model of the used car market outlined above. It will become apparent, however, that there is nothing special about the used car market. Similar conclusions would arise in many markets where adverse selection may be a serious problem. In Section 2, I describe the model in detail and apply the conventional analysis of the market with adverse selection in Section 3. The equilibrium in the conventional analysis is interpreted as the outcome generated by an auctioneer who has instructions to set a price which equates supply and demand. I argue that, aside from the possibility of perverse income effects, the presence of adverse selection may lead to multiple equilibria. Furthermore, these equilibria can always be ranked according to the Pareto criterion in order of ascending price.

The remainder of the paper is devoted to examining how the nature of the equilibrium is affected when the agents on one side of the market are explicitly assigned the role of setting the prices. In Section 4 I examine the nature of the equilibrium under the explicit assumption that buyers are the price setters. Each buyer must announce a fixed price at which he is willing to purchase one car. Sellers are then free to submit bids to as many buyers as they wish and to sell at the highest bid accepted. If there is excess supply at any price, it is rationed at random. With this convention specified, I ask the following questions: Will all buyers necessarily announce the same price and do their prices necessarily equate supply and demand? The answer to both of these questions is no.

The basic argument is as follows. Suppose some buyer announces a higher price. Besides attracting all sellers who submitted bids at lower prices, the buyer will also attract some new sellers. Even though the buyer may not be able to identify the new sellers, the addition of their cars may nevertheless increase the average quality of cars sufficiently so that the expected benefit to the buyer has increased. In such a case he has no incentive to lower his price, even though he faces excess supply. However, not all buyers need have such an incentive. Those buyers who value increases in quality less may prefer to announce a lower price. They will purchase cars from those sellers who were unable to sell at the higher price. Such a market will be characterized by a distribution of prices with excess supply at all but the lowest price.

In Section 5 I focus on the opposite convention. Sellers announce a fixed price and each buyer submits a bid to the sellers announcing the price at which he wishes to purchase a car. Again, I assume that if there is excess supply at any price, it is rationed at random. Unlike the case in which buyers set the price, however, a seller is no longer free to change his price if he is unable to find a buyer. The price that a seller chooses, therefore, typically depends upon both his expected probability of making a sale at each price and the value he attaches to his car. The *actual* probability of selling a car at any price, however, depends on how many buyers submit bids at that price, which in turn depends upon their expectations about how quality is related to price. Consequently,

depending on the expectations of both buyers and sellers, the equilibrium can take on a number of different forms. If the expectations of sellers are such that owners of higher quality cars choose to announce higher prices and buyers are sufficiently optimistic about the relation between quality and price, the equilibrium may be characterized by a distribution of prices. If buyers are more pessimistic about the relation between price and quality, however, all sellers will be induced to announce the same price. In any case, there is always a family of single-price equilibria.

The final sections of the paper are concerned with the robustness of the equilibrium under different price-setting conventions and a comparison of some of their welfare implications.

2. The model

■ The model consists of a fixed number of cars of varying quality and a set of agents. Each car is assigned a quality index, $q > 0$. Each agent has a von Neumann-Morgenstern utility function, characterized by a parameter, t , which defines his relative valuation of a car of quality q to consumption, c :

$$u(c, q; t) = c + tq.$$

If an agent does not own a car, q may be set equal to 0. Assume the set of agents may be partitioned into two subsets: those that initially own one car and those that initially own none.

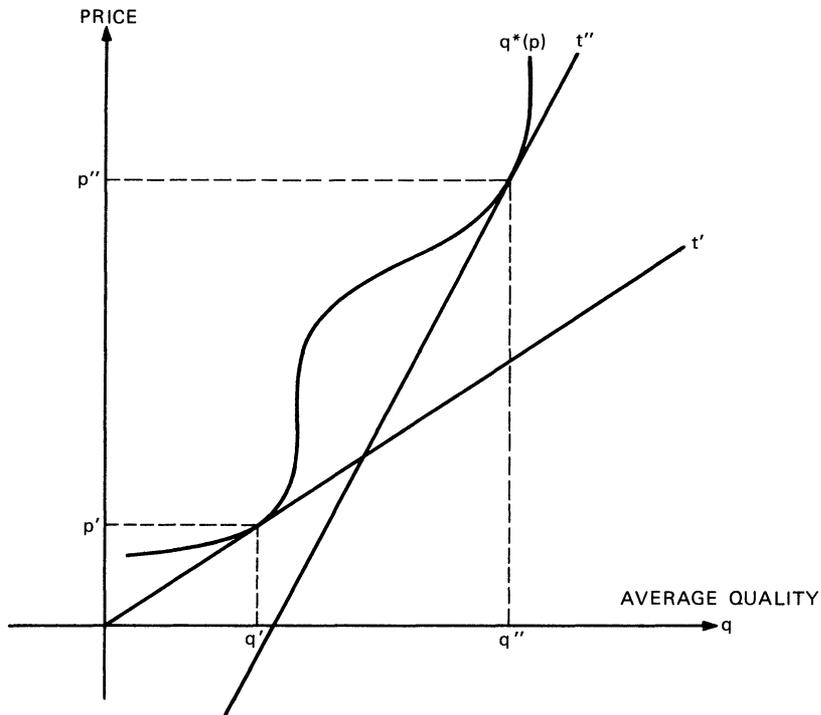
□ **Nonowners.** Nonowners may be characterized by their utility index, t , which I assume is distributed according to a continuously differentiable and strictly positive density, $h(t)$, defined on $[t_0, t_1]$, where $t_0 > 0$. Let $H(t) = \int_{t_0}^t h(x)dx$. I shall sometimes refer to $H(t)$ as the number of buyers with utility index less than t .¹

The special form of the utility function implies a number of properties which will be exploited in the remainder of the paper. First, since utility is linear in q , buyers are risk-neutral with respect to quality. When faced with a distribution of cars from which the buyer may choose one, his expected utility depends only on the expected quality of cars in that distribution. Second, for any given quality of car, the higher the utility index, t , the higher the reservation value the buyer assigns to the car. This implies that the set of buyers who enter the used car market can always be characterized by a critical utility index t^* . Nonowners enter the market if and only if $t \geq t^*$. Third, the higher the utility index t , the higher a buyer's marginal rate of substitution between quality and consumption of other goods. This implies that when buyers are faced with a distribution of prices where average quality depends on price, buyers with higher utility indices will choose to purchase cars at prices at least as high as the prices chosen by buyers with lower utility indices. Proposition 1 summarizes these last two properties.

Proposition 1: Given an expected quality function defined over a set of prices P , if buyer t' chooses to purchase at price p , then any buyer $t'' > t'$ will choose to purchase at a price at least as high as p . In particular, if t' chooses to enter the market at price p , any $t'' > t'$ will also choose to enter the market.

¹ In particular, I do not require that $H(t_1)$ equal 1.

FIGURE 1
BUYERS' PRICE IS AN INCREASING FUNCTION OF t



The argument is illustrated in Figure 1. Price is measured on the vertical axis and average quality on the horizontal axis. The curve labelled $q^*(p)$ represents an arbitrary schedule of the average quality of cars which can be purchased at each price. A typical indifference curve for a buyer with utility index t' is represented by a straight line with slope t' . Increasing q and lowering p increase utility. Therefore, buyer t' maximizes utility by purchasing at price p' and receiving expected quality q' . Note that he is indifferent to buying at p' or not purchasing any car at all. Buyers with utility index $t < t'$ will strictly prefer to remain out of the market. As the utility index rises to t'' , the slope of the indifference curve increases; as a consequence, the optimal price for buyer t'' must be at least as high as p' . In Figure 1, the optimal price for t'' is $p'' > p'$. Note that even if buyer t'' were permitted only to purchase at price p' , he would still enter the market.

□ **Owners.** All agents in the ownership set are assumed to have the same utility index, t . They differ only in the quality of cars they own, which I again assume is distributed according to a continuously differentiable and strictly positive density, $f(q)$, defined on $[q_0, q_1]$ with $q_0 > 0$. Let $F(q) = \int_0^q f(x)dx$. Then $F(q)$ may be interpreted as the number of sellers with cars of quality less than q .

To ensure that there is some incentive for trade to take place, it will also be useful to assume that some nonowners have a stonger preference for a car of given quality than do the owners. Also, the argument in Section 5 will be simplified if some nonowners have a lower utility index than the sellers. Therefore, I shall make the additional assumption that $t_0 < t < t_1$.

The model outlined above is admittedly very special. All agents are assumed to use the same measure of quality, and utility is assumed to be linear in quality. Furthermore, all sellers have the same utility function. I have chosen to restrict my analysis to this specific model only because most of the computations are substantially simplified and the economic principles behind my results become more apparent. For most of the results presented below, these assumptions can be weakened considerably.

3. Walrasian equilibrium with adverse selection—buyers and sellers are price takers

■ In the remainder of the paper, I assume that buyers cannot directly determine the quality of car they purchase. However, they can observe the average quality of all the cars sold and the prices at which they are sold.

Suppose price is set by an auctioneer to equate supply and demand. If the auctioneer is unable to discriminate among cars of different quality, all cars will sell for the same price. The supply of cars at each price is then equal to the number of owners for whom the price exceeds the benefit of owning their cars. The demand at each price is equal to the number of buyers for whom the expected benefit from buying a car exceeds the price. This number depends not only on the price but also on the average quality of the cars supplied by sellers at that price. A *Walrasian equilibrium* is defined as the price at which the quantity of cars supplied is equal to the quantity of cars demanded.

The value an owner assigns to owning a car of quality q is equal to tq . Therefore, at price p , an owner will sell his car if and only if $p \geq tq$. This implies that the Walrasian supply at price p , $S(p)$, is equal to the number of cars for which $q \leq p/t$:

$$S(p) = \begin{cases} \int_{q_0}^{p/t} f(q) dq & \text{for } p > tq_0, \\ 0 & \text{otherwise.} \end{cases}$$

The average quality of the cars offered for sale at price p is then given by

$$q^a(p) = \begin{cases} \int_{q_0}^{p/t} qf(q) dq / S(p) & \text{for } p > tq_0, \\ q_0 & \text{for } p = tq_0. \end{cases}$$

I shall refer to q^a as the *Walrasian average quality* function. Since no sellers offer their cars at prices less than tq_0 , q^a is defined only at prices at or above tq_0 .

The assumption that all sellers have identical utility functions implies a strong relation between the average quality and the price. Since the quality of the marginal seller's car exceeds the quality of any other car offered for sale, the quality of the marginal car must be higher than the average. Therefore, increases in the price must increase average quality.²

Assuming that nonowners can observe only the average quality of the cars

² If owners differ not only by the quality of their cars, but also by the value of their utility indices, the same result would still be obtained if the distribution of sellers by utility indices is independent of the distribution of the quality of their cars. Without the independence assumption, however, it is possible that the opposite relation could result.

sold at each price, the expected benefit of buying a car at price p for a buyer with utility index t is $tq^a(p)$. Since buyers will only enter the market if their expected benefit exceeds the price, the Walrasian demand at price p , $D(p)$, is equal to the number of buyers with utility index $t \geq p/q^a(p)$ ³:

$$D(p) = \begin{cases} \int_{p/q^a(p)}^{t_1} h(t)dt & \text{for } p < t_1q^a(p), \\ 0 & \text{otherwise.} \end{cases}$$

Assuming that the auctioneer is instructed to set a price which equates supply and demand, we may then define:

A *Walrasian equilibrium* (W.E.) is a price p^e such that $D(p^e) - S(p^e) = 0$.

This is essentially an outline of the analysis presented by Akerlof (1970). By setting $q_0 = 0$ and assuming that the distribution of q was uniform, he was able to construct an example where the demand completely collapsed and the only equilibrium price was zero. This need not always be the result. In fact, the assumptions of Section 2 (in particular, the requirements that $t_1 > t$ and $q_0 > 0$) imply that there is always an equilibrium with positive price and a positive level of trade.⁴

The equilibrium price need not be unique. As usual, supply is an increasing function of price (strictly increasing for $tq_0 < p < tq_1$). The demand function, however, is not necessarily downward sloping. The reason has nothing to do with perverse income effects. Unlike the situation in the conventional theory of demand, in this model the nature of the product depends upon the price. If as the price rises, the increase in average quality is sufficiently high, the benefit to the marginal buyer may rise by more than the price. At such prices demand will have a positive slope. Since the utility index of the marginal demander is defined by $t = p/q^a(p)$, the necessary and sufficient conditions for an upward sloping demand function reduce to

$$\frac{p}{q^a(p)} \frac{dq^a(p)}{dp} > 1. \tag{1}$$

If, in addition to (1) being satisfied, the density of buyers at the marginal value of t is sufficiently high, the demand curve may become even flatter than the supply curve. In this case, there may be multiple equilibria.

An example is illustrated in Figure 2. The supply function is represented in the northeast quadrant. From the definition of supply, $S(p) = 0$ for $p < tq_0$ and $S(p) = F(q_1)$ for $p > tq_1$. The relation between price and the ratio of price to the average quality of cars is represented in the northwest quadrant. At $p = tq_0$, average quality is q_0 , and therefore $p/q^a(p) = t$. As price rises, the

³ I am implicitly assuming that buyers may purchase at most one car. Also they are not permitted to buy and then resell a car in the same market.

⁴ At price tq_0 , supply is zero and demand is positive. At $p = tq_1$, supply is positive but, since $q_1 > q^a(p)$, demand must be zero. The existence of a positive price equilibrium then follows from the continuity of S and D .

Whether or not any price below tq_0 can be an equilibrium is somewhat problematic. Since no sellers offer their cars at that price, average quality and, therefore, demand are not strictly defined. If we adopt the convention that zero supply implies zero average quality, then $p = 0$ is always an equilibrium price.

is a single price equilibrium. By repeating this procedure at all prices, one may construct the entire demand curve. Since $p/q^a(p)$ was assumed to be decreasing over a range of prices, the demand function is increasing over that range. Furthermore, the density of the distribution of buyers is sufficiently high at those prices that demand also intersects supply at p_2 and p_3 . Therefore, p_1 , p_2 , and p_3 are all Walrasian equilibria.

In those cases where there are multiple equilibria, the equilibria can be ranked according to the Pareto criterion with higher price equilibria generating higher welfare for all active agents. For the sellers, the ranking is obvious. As long as he can sell his car with certainty, any seller would prefer a higher to a lower price. For the buyers, the result follows from the assumption, implicit in the form of the utility function, that the lower is a nonowner's reservation price, the higher is his marginal rate of substitution of quality for price. For two distinct prices both to be equilibria, demand must be higher at the higher price (otherwise, demand could not equal supply at both prices). This implies that there are some nonowners who choose not to purchase at the lower price, but who will purchase at the higher price. These agents are evidently better off at the higher price. But since these agents must have a lower utility index t than any agent who buys at the low price, and since t is a measure of the agents' marginal rate of substitution between quality and price, it follows that quality must have risen sufficiently so that *every* buyer who is willing to purchase at the lower price prefers to purchase at the higher price. Thus, the higher price is Pareto preferred.

Wilson (1978) develops this argument in more detail. There I also argue that in some cases, even prices which generate excess supply may be Pareto preferred to any market clearing price. I shall return to this point briefly in Section 6.

4. Buyers set the price

■ In this section I retain the assumption that buyers can observe only the average quality of the cars actually sold at each price, but suppose that prices are explicitly set by the buyers. The market is assumed to operate according to the following rules. Each buyer may announce at most one price at which he is willing to purchase a car. Once the announcements are made, each seller may offer his car to any buyers he wishes. If there is excess supply at any price, sellers are rationed at random. Those unable to sell at the higher prices may offer their cars to buyers with a lower price. This process continues until each owner either sells his car or decides to keep it.

The behavior of sellers in this kind of market is relatively uncomplicated. They try to sell at the highest price possible. If a seller is unable to sell at any price above the reservation value of his car, he leaves the market.

The problem facing buyers is more complicated. The price a buyer announces depends on both the quality of car he expects to receive at each price and the likelihood of finding a seller at that price. Depending on their preferences, some buyers may prefer to announce prices which are higher than necessary to attract sellers, if the increase in expected quality of cars outweighs the higher price. Other buyers may choose to announce no price at all, because, at any price they are willing to announce, they do not expect any owner to sell.

The market is in equilibrium when the expectations of the buyers are

correct. Assuming that buyers correctly anticipate both the set of prices at which a car can be purchased and the average quality of cars forthcoming at each price, a unique equilibrium will exist. In this equilibrium, either all buyers will announce the highest Walrasian equilibrium price or else buyers will announce a distribution of prices extending above and below this price.

□ **Buyers' expectations and demand.** If a nonowner wishes to purchase a car, he must announce a price at which he will accept an offer to sell. Once announced, this price cannot be altered. The price he chooses will depend upon (1) the average quality of the cars he expects to be offered for sale at each price and (2) the likelihood that some seller will sell at the price he selects. Assuming that sellers can search costlessly for the most favorable price, buyers may expect an offer at price p only if there is an excess supply of cars at all prices higher than p . Therefore, I shall assume that buyers make a point estimate of a critical *cutoff price* p . At prices at or above p , their subjective probability of making a purchase is one; at prices below p , the probability is zero.⁶ The average quality of cars buyers expect to be forthcoming at each price may be summarized by an expected quality function, $q^e(\cdot)$, which I assume to be a continuous function of price. Both p and q^e are identical for all buyers.⁷ Given these assumptions, there is an optimal price announcement for each buyer.

□ **Sellers' supply and the average quality functions.** There are two extreme assumptions one might make about the opportunities of sellers who are faced with a given announcement of prices. If the costs of search are high, we might assume that owners may offer their cars to buyers at only one price. If such costs are negligible, however, it may be more appropriate to assume that each seller may offer his car to as many buyers as he wishes until the car is sold. The first assumption leads to an analysis very similar to that in the next section, where I assume sellers set the price.⁸ Therefore, I shall confine my attention in this section to the second assumption. Sellers may offer their cars to as many buyers as quote a price; each seller then sells his car at the highest price at which his offer is accepted. If the supply of cars exceeds the demand at any price, sales at that price are distributed at random.

Given these assumptions, there are two reasons why the quantity of cars supplied will be an increasing function of the price. First, the higher the price, the greater the percentage of offers that are not withdrawn because the seller has found a higher price. Second, the higher the price, the greater the number of sellers with lower reservation values, and consequently, the greater the number of offers to sell. The net effect on supply can best be illustrated in a simple example.

Suppose buyers announce two prices: $p_1 < p_2$ and let $b(p_i)$ be the number of buyers who announce price p_i . At any price p , a seller with a car of quality q will offer his car for sale only if: (1) $p > tq$; and (2) no buyer has agreed to

⁶ One might argue that the probability of receiving an offer at p could lie anywhere between zero and one. However, the assumptions below imply that, in equilibrium, buyers will always be able to buy a car at the cutoff price. Therefore, it is not really an additional restriction. I assume it from the outset only to ensure that the buyers' optimal strategies are well defined (see footnote 7).

⁷ This assumption also is not restrictive, since it will have to be satisfied in equilibrium anyway.

⁸ This is the approach taken by Stiglitz (unpublished) in a similar analysis of labor markets.

buy his car at a higher price. Recall that $S(p) = \int_{q_0}^{p/t} f(q) dq$. Therefore, at the highest price, p_2 , the excess supply may be written:

$$E(p_2) = S(p_2) - b(p_2) = S(p_2)[1 - b(p_2)/S(p_2)]. \tag{2}$$

Assuming $E(p_2) > 0$, the excess supply at price p_1 is then equal to the number of sellers unable to sell their cars at the higher price, $E(p_2)$, times the proportion of sellers who remain in the market at the lower price, $S(p_1)/S(p_2)$, minus the demand at p_1 , $b(p_1)$. Therefore, (2) implies

$$\begin{aligned} E(p_1) &= [S(p_1)/S(p_2)]E(p_2) - b(p_1) \\ &= S(p_1)[1 - b(p_2)/S(p_2) - b(p_1)/S(p_1)]. \end{aligned}$$

This argument can be extended by induction as long as excess supply remains positive. The result is an expression for $E(p_1)$ when any finite number of higher prices is announced by sellers:

$$E(p_1) = S(p_1)[1 - \sum_{i=1}^n (b(p_i)/S(p_i))]. \tag{3}$$

The excess supply function for a continuum of prices has a form similar to (3). For each cutoff price p and expectation function q^e , let $B(p; p, q^e)$ represent the number of buyers who announce a price less than p , and let $B_+(p; p, q^e)$ be the number of buyers who announce a price less than or equal to p . Then, if buyers have cutoff price p and expectation function q^e , the excess supply of cars at each price is:

$$E(p; p, q^e) = \begin{cases} S(p) \left[1 - \int_p^\infty \frac{dB(x; p, q^e)}{S(x)} \right] & \text{if } E(p'; p, q^e) > 0 \\ \text{for all } p' > p \text{ and } p > tq_0, \\ B(p; p, q^e) - B_+(p; p, q^e), & \text{otherwise.} \end{cases}$$

Because sellers may costlessly search for the highest possible price and because sellers are selected at random whenever there is excess supply, the average quality of cars offered for sale at each price continues to be defined by the Walrasian average function defined in the previous section. The average quality of cars offered at each price equals the average quality of the cars owned by sellers with reservation values lower than that price.

□ **Equilibrium.** Suppose that buyers are able to observe the set of prices at which cars are actually supplied and the average quality of the cars supplied at each of these prices. Assuming the buyers are able to adjust their expectations in the light of these observations, at a minimum, equilibrium should be characterized by the following two conditions. First, at every price at which a sale occurs, the buyers' expected quality function should equal the actual quality function. Second, at every price at or above the estimated cutoff price p , the excess supply of cars must be nonnegative. If both of these conditions are satisfied, the expectations of buyers will not be explicitly invalidated.

If this were all we were to require of an equilibrium, then any Walrasian equilibrium could be sustained. For instance, suppose buyers were unaware that a seller's reservation price is correlated with the quality of car he owns. In this case, the quality of car they expect to be offered would be independent

of the price. Competition among buyers would then force the price to a point where the supply of cars equals the number of buyers willing to announce that price. Assuming that buyers adjusted their expected cutoff price to this level, and adjusted their expected quality function appropriately, this expectation would not be unsubstantiated and the market would be in equilibrium.

In addition, the equilibrium price would not necessarily have to equate supply and demand. If buyers were aware that average quality increases with price, but were overly pessimistic about the decrease in quality resulting from a lower price, other single price equilibria could also be sustained. At these prices there would be an excess supply of cars; however, given the expectations about the relation between quality and price, no buyer would have an incentive to lower his price. An example of a one parameter family of such equilibria is given in an unpublished version of this paper (Wilson, 1977).

One might reasonably argue, however, that a concept of equilibrium which requires that expectations be confirmed only on the set of announced prices is not sufficiently stable. If the expected quality function does not agree with the actual quality function at prices outside the announced set or if the estimated cutoff price is too high, such equilibria may not be robust to experimentation by the price setting buyers. For instance, in the example just cited, if one buyer experiments by announcing a price either higher or lower than the equilibrium price, the buyers' quality function will be invalidated and they will revise their expectations accordingly. Those equilibria most likely to be robust, therefore, are those for which, given the true average quality function and the true cutoff price, no buyer would choose a different price. This restriction leads to the following definition:

A pair (p, q^e) is a *buyers' equilibrium* (B.E.) if:

$$(a) \quad q^e(p) = q^a(p) \quad \text{for all } p \geq tq_0;$$

$$(b) \quad E(p; p, q^e) = 0 \quad \text{if and only if } p \leq p$$

The existence of a buyers' equilibrium is established in an appendix that is available from the author on request. There I also show that the equilibrium is unique. Since the expected quality function must equal the Walrasian average quality function at all prices, there can be more than one equilibrium only if there is more than one equilibrium cutoff price. But this is not possible, since a decrease in the cutoff price can only increase the number of buyers and decrease the number of sellers who are able to sell a car before leaving the market.

□ **Properties of the buyers' equilibrium.** We turn now to a more detailed description of the buyers' equilibrium. Depending on the shape of the average quality function and the distribution of preferences, there are essentially two forms the equilibrium may take. Under some conditions, every buyer will announce the highest Walrasian equilibrium price. In all other cases, different buyers will announce distinct prices resulting in a distribution around the highest Walrasian equilibrium price.

To see when a Walrasian equilibrium will emerge, let p^* be a W.E. price and consider a buyer with the highest utility index t_1 . If he prefers to buy a car at p^* rather than at any higher price, then Proposition 1 implies that all buyers prefer p^* to any higher price. Therefore, if we set the cutoff price equal to p^* ,

all buyers who enter the market will announce p^* . But by definition, p^* equates supply and demand; therefore, excess supply at p^* must be zero and (p^*, q^a) is the B.E. Note, however, that this condition can never be satisfied unless p^* is the highest W.E. price. If p^{**} is a higher W.E. price, then since supply is an increasing function of price, it follows that there exist some nonowners who will purchase at p^{**} but not at p^* (otherwise, supply could not equal demand at both prices). Proposition 1 then implies that every buyer would prefer to announce p^{**} . These results are summarized in Proposition 2.

Proposition 2: Let p^* be the highest W.E. price and let (p, q^a) be the B.E. If, given quality function q^a , buyer t_1 prefers to purchase at p^* rather than at any higher price, then $p = p^*$, and all active buyers announce p^* .

The argument behind Proposition 2 is quite general. So long as no buyer prefers to buy at a higher price, the market equilibrium is no different from the one resulting when an auctioneer equates supply and demand. When some buyers do prefer to announce a higher price, however, this conclusion must change.

Proposition 3: Let p^* be the highest W.E. price and let (\bar{p}, q^a) be the B.E. If buyer t_1 prefers to purchase at a price higher than p^* , then p must be less than p^* .

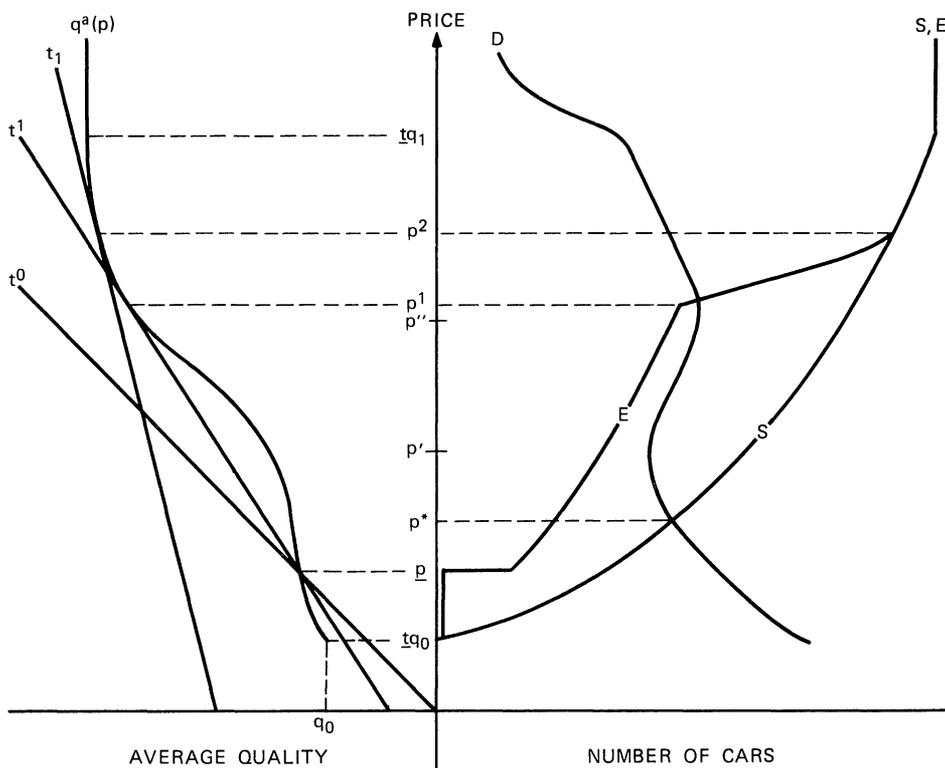
A formal statement and proof of this result are given in the unpublished appendix. Here I shall present only a heuristic argument. Suppose first that not all buyers who demand a car at p^* prefer to announce a higher price. Then from Proposition 1 and the definition of $D(p)$, any buyer who announces a price greater than p^* would also demand a car if faced with the price p^* . Therefore, the number of buyers who announce a price greater than or equal to p^* cannot exceed $D(p^*)$. But, buyers who announce prices greater than p^* attract some sellers who would not supply a car at price p^* . Hence, at least some of the sales at any announced price above p^* will be made by sellers who would not offer their cars at p^* . As a result, if any price above p^* is announced, the excess supply of cars at p^* must exceed the excess supply—namely, zero—one would find at p^* if only p^* were announced.

Next suppose that all buyers who are willing to purchase at p^* also prefer to announce a price which is greater than p^* . By the definition of p^* , at any price $p > p^*$, $S(p) - D(p) > 0$. Consequently, some owners with a reservation price less than p^* will be unable to sell at a price greater than p^* . Again, therefore, excess supply must be positive at p^* . Since, by definition of a buyers' equilibrium, excess supply must be zero at the cutoff price, we have established that in either case, the cutoff price must lie below p^* .

Although the proof of Proposition 3 relies on the form of the buyers' utility function, it does not depend on any special assumptions about the preferences of sellers. Because all sellers have the same utility function, however, we can also establish that at least one buyer must announce the cutoff price. To see this, suppose the contrary. Then there must be excess supply at the lowest price actually announced; otherwise, it would be the cutoff price. Therefore, since there is no demand between the lowest announced price and tq_0 , to reduce excess supply to zero, the cutoff price must equal tq_0 . At price tq_0 , however, only cars of quality q_0 are supplied. Therefore, all buyers with $t \geq t$ would be willing to buy at that price. At any price above tq_0 , however, the

FIGURE 3

THE EQUILIBRIUM EXCESS SUPPLY FUNCTION



quality of the marginal car exceeds the average. Since no owner will sell if $p \leq tq$, this implies that $p/q^a(p) > t$ for all $p > tq_0$. From this it follows that the utility index of the marginal buyer must be strictly larger than t . Consequently, if excess supply were strictly positive at the lowest price actually announced, some buyers with t slightly larger than t would enter the market and announce a price near tq_0 . Since this would attract some of the sellers unable to sell at a higher price, the original set of announced prices cannot be consistent with equilibrium.

An example of the relation between the equilibrium excess demand function, E , and the Walrasian supply and demand functions is presented in Figure 3. Note that S and D intersect only once at p^* , so that the Walrasian equilibrium is unique. To make the example more interesting, however, I have constructed the demand curve to be upward sloping from p' to p'' . This reflects the fact that the elasticity of the average quality function q^a , drawn on the left-hand side of the price axis, is greater than one over this range. Now suppose we have already determined that p is the equilibrium cutoff price. From Propositions 2 and 3 we know that $p \leq p^*$. The price announced by each buyer is determined by the point on the average quality function at or above p which reaches his highest indifference curve. This will be at p or tq_1 or a point of tangency. Buyer t^1 announces p^2 . Buyer t^1 is just indifferent between p^1 and p . Buyer t^0 is the lowest buyer who enters the market; he also announces p . From this information we can get a rough idea of the shape of the excess supply function.

At prices at or above p^2 , there is no demand; therefore, the E function and S functions are identical. From p^2 to p^1 additional buyers continuously enter the market; as a consequence, the excess supply function becomes flatter than S . From p to p^1 , there are no buyers; therefore, the decline in E as p falls is proportional to S . At p , just enough buyers enter the market to eliminate the excess supply.

There are two points to note about this example. First, no buyer announces a price in the interval where the demand curve has a positive slope. Given the form of the utility function, this property is quite general. The reader may readily verify that any buyer willing to purchase at a price where the demand curve is rising would prefer to purchase at an even higher price.

Secondly, note that there is a mass point of buyers who announce p , while buyers are distributed continuously at all the higher prices. In general, this is the form we should expect the equilibrium to take. Buyers who announce a price greater than p are attaining an interior maximum. Given our assumptions on the distribution of sellers, $q^a(p)$ is continuously differentiable from tq_0 to tq_1 . Therefore, any price greater than p and less than tq_1 can be optimal for at most one utility index t . In contrast, buyers who announce price p are not attaining an interior maximum. Some of them may prefer an even lower price. However, if these buyers were to announce their optimal prices and extend the distribution below p , excess supply would have to be negative over some interval of these prices. The buyers would then be forced to raise their prices to attract sellers, thereby forcing other buyers to raise their prices until all were announcing the cutoff price p . I have already argued that at least one buyer must announce the cutoff price. If the demand curve is negatively sloped at all prices below p^* , one may also show that p must be announced by a mass point of buyers.⁹

5. Sellers set the price

■ In this section the roles of the price makers and the price takers are reversed. Retaining the assumption that buyers can observe only the average quality of the cars offered at each price, I investigate the nature of the equilibrium when the sellers explicitly set the price. The market works as follows. Each seller has the option of announcing a fixed price or staying out of the market. Given the set of announced prices, each buyer then decides at which price, if any, he will purchase a car. If there is an excess supply of cars at any price, it is assumed that the probability of selling a car at that price is equal to the ratio of demand to supply.

Because their prices must remain fixed even after buyers have submitted their bids, the prices sellers choose will depend on the expected probability of making a sale at each price. Given the set of prices announced by the sellers, buyers must then form an expectation about the quality of cars at each price. Having formed their expectations, they purchase at the price which maximizes their expected utility. Equilibrium is defined to be an expected probability function for sellers and an expected quality function for buyers such that the behavior resulting from these expectations generates a realized probability

⁹ A mass point is also possible at tq_1 since q^a is not differentiable at this price. In fact, unless there is excess demand at this price, no buyer will ever announce a price greater than tq_1 since higher prices no longer indicate higher quality.

function and a realized quality function which are consistent with the original expectation functions.

Perhaps the most important difference between the buyers' equilibrium and the sellers' equilibrium is that when the sellers set the prices, the equilibrium average quality of cars received by buyers at each price is no longer independent of the prices at which other buyers choose to purchase. Consequently, the equilibrium is also no longer unique. Since the prices announced by sellers depend on their expectations of selling a car at each price, and the response of buyers depends on their expectations of average quality at each price, many different patterns of expectations can be self-confirming. Nor does this indeterminacy in the nature of equilibrium stem from the lack of detailed knowledge of the structure of the market on the part of buyers or sellers. Even if all agents know the distribution of the preferences of buyers and the quality of cars of sellers, the equilibrium may still take a number of different forms.

A complete analysis of the equilibrium under this price-setting convention is somewhat lengthy and requires considerable attention to technical detail. It is contained in the unpublished appendix. To focus on the essential features, therefore, I shall confine my attention to two extreme cases. For any distribution of buyers and sellers satisfying the assumptions of Section 2, there is a family of equilibria where all sellers announce the same price. There is also a family of equilibria where all sellers announce distinct prices. Over the interval of prices which are announced, quality increases with the price, while the probability of making a sale decreases with the price.

□ **Equilibrium.** It is convenient to think of the market as working in two steps: the sellers announce prices and then the buyers respond. The expectations of the sellers are summarized by a function, $\pi^e(p)$, which defines the expected probability of making a sale at each price. Given this function, each seller announces the price that maximizes his expected utility. (If $\pi^e(p) = 0$ for all prices above his reservation price, he announces no price at all.) The combined actions of all sellers then generate a realized quality function $q^*(p; \pi^e)$ equal to the average quality of cars offered for sale at each price p , and a supply function $s(p; \pi^e)$ equal to the number of sellers who announce each price.

Given the set of prices announced by sellers, $P(\pi^e)$, buyers then formulate their estimate of the quality of car offered at each price. This is summarized by an expected quality function $q^e(p; P(\pi^e))$, which defines the expected quality of cars at each price in $P(\pi^e)$. Each buyer then purchases at the price that maximizes his expected utility. This in turn generates a demand function $b(p; P(\pi^e), q^e)$ equal to the number of buyers at each price. Assuming that $b(p; P(\pi^e), q^e) \leq s(p; \pi^e)$, the ratio $b(p; P(\pi^e), q^e)/s(p; \pi^e)$ then defines the probability of making a sale at each price, $\pi^*(p; \pi^e, q^e)$.

I shall say that the market is in a *sellers' equilibrium* when the expectations of neither the buyers nor the sellers are invalidated. That is, on the set of prices actually announced, the realized probability function equals the sellers' expected probability function, and the realized quality function equals the buyers' expected quality function. Formally,

A pair of expectation functions (q^e, π^e) is a *sellers' equilibrium* (S.E.) if for all $p \in P(\pi^e)$:

- (i) $q^*(p; \pi^e) = q^e(p; P(\pi^e))$;
 (ii) $\pi^*(p; \pi^e, q^e) = \pi^e(p) \leq 1$.¹⁰

In the remainder of the section, it will be convenient to delete the π^e and q^e arguments in the q^* and π^* functions.

□ **Single price equilibria.** Note first that this definition of equilibrium is consistent with the Walrasian equilibrium. In fact, with an appropriate choice of the sellers' expected probability function, it can be made consistent with the announcement of any single price at which excess supply is nonnegative.

To see this, choose any price p_0 at which the Walrasian supply $S(p_0)$ is greater than or equal to the Walrasian demand $D(p_0)$. Now define the expected probability function $\pi^e(p)$ to equal $D(p_0)/S(p_0)$ for $p \leq p_0$ and 0 for $p > p_0$. Then any seller who enters the market will choose to announce p_0 . Therefore, the set of announced prices, $P(\pi^e)$, is simply $\{p_0\}$. Furthermore, since the same sellers enter the market as would enter if an auctioneer were announcing p_0 , we also have $s(p_0) = S(p_0)$ and $q^*(p_0) = q^a(p_0)$. Given the actions of sellers, buyers only need to formulate their expected quality function at p_0 . If we then suppose that $q^e(p) = q^a(p)$, not only will their expectations be confirmed, but utility maximization will imply that the number of buyers, $b(p_0)$, will equal the Walrasian demand at that price, $D(p_0)$. Therefore, we have $\pi^*(p_0) = D(p_0)/S(p_0) = \pi^e(p_0)$, so that the sellers' expectations are confirmed on the set of announced prices. From this, it follows that (q^e, π^e) is a *sellers' equilibrium*.

This example also illustrates the major weakness of any purely expectational equilibrium concept. Since it puts no restrictions on expectations outside the set of prices actually observed, it allows for the possibility that agents may persistently behave very inefficiently simply because their expectations outside the set of announced prices are never tested. I shall later return to this issue. For the present, however, we shall be concerned with finding those patterns of expectations on the part of buyers and sellers that are mutually consistent.

□ **Discriminating sellers' equilibrium.** The polar opposite of an equilibrium in which all sellers announce the same price is one in which sellers of different quality cars announce distinct prices. I shall call such an equilibrium a *discriminating sellers' equilibrium* (D.S.E.). Equilibria of this form are possible only because of the relationship between the quality of a seller's car and his tradeoff between price and the probability of making a sale.

Because owners of higher quality cars have higher reservation values, they have less to lose if they are unable to make a sale at any given price. Consequently, for a given increase in the price, owners of higher quality cars are willing to accept a larger decrease in the probability of making a sale. This

¹⁰ Note that the definition of equilibrium explicitly excludes the possibility of excess demand at any price. This restriction removes the necessity of defining how buyers will respond if they are unable to purchase at their most preferred price. It is possible to formulate the equilibrium concept to include the possibility of excess demand. However, it is easy to show that excess demand could only appear at the lowest price announced by any seller. Furthermore, using the arguments to be presented in a later subsection, one can argue that experimentation by sellers would eliminate any such equilibria.

implies that the expected quality function can be adjusted so that sellers of higher quality cars choose to announce higher prices, thereby generating an upward sloping quality function. Depending upon the shape of this function, some buyers may then have an incentive to purchase at prices above the minimum and thus generate a demand function which confirms the sellers' expected probability function. In fact, the expected quality and expected probability functions can be appropriately adjusted to yield a one-parameter family of these equilibria.

The problem of constructing a D.S.E. may be reduced to finding a solution to a system of three differential equations. One equation defines the distribution of buyers, one defines the distribution of sellers and hence the realized quality function, and the last equation defines the realized probability function.

Consider the sellers first. Given an expected probability function equal to the realized probability function π^* , a seller with car quality q will choose the price that maximizes $\pi^*(p)(p - tq)$. Since, by definition, a seller with a car of quality $q^*(p)$ chooses p , the first-order conditions for an interior maximum imply

$$\pi^{*'}(p)[p - tq^*(p)] + \pi^*(p) = 0. \quad (3)$$

Turning to the buyers, Proposition 1 implies that for any expectation function, buyers with higher utility indices will always purchase at higher prices. Let $t^*(p)$ be the utility index of a buyer who purchases at price p . Since the objective of each buyer is to choose the price that maximizes $tq^*(p) - p$, the first-order conditions for a maximum imply

$$t^*(p)q^{*'}(p) - 1 = 0. \quad (4)$$

When both the buyers and sellers are distributed continuously over an interval of prices, the realized probability function must be defined with some care. Because the set of agents at any price is negligible, it makes no sense to talk about the ratio of buyers to sellers at any single price. Rather, we must consider the number of buyers and sellers in arbitrarily small intervals around each price. Since the number of buyers per unit interval is $h(t^*(p))t^{*'}(p)$ and the number of sellers is $f(q^*(p))q^{*'}(p)$, the realized probability of selling a car at price p may, therefore, be defined as

$$\pi^*(p) = \frac{h(t^*(p))t^{*'}(p)}{f(q^*(p))q^{*'}(p)}. \quad (5)$$

Equations (3)–(5) may then be rewritten to define a system of three first-order differential equations:

$$\pi^{*'}(p) = - \frac{\pi^*(p)}{p - tq^*(p)} \quad (6)$$

$$q^{*'}(p) = \frac{1}{t^*(p)} \quad (7)$$

$$t^{*'}(p) = \frac{\pi^*(p)f(q^*(p))q^{*'}(p)}{h(t^*(p))} = \frac{\pi^*(p)f(q^*(p))}{t^*(p)h(t^*(p))}. \quad (8)$$

Utility maximization also implies some restrictions on the values of q^* , t^* , and π^* at the end points of any distribution. Suppose the subset of announced prices is the interval $[p_0, p_1]$. Then since equation (7) implies that $q^*(p)$ is

increasing in p , and since no seller announces a price lower than his reservation value, we have

$$q^*(p_0) = q_0; \quad p_0 - tq_0 > 0. \tag{9}$$

Since $t^*(p)$ is also increasing in p , buyer $t^*(p_0)$ must be just indifferent to entering the market; otherwise buying at p_0 would be optimal for an interval of buyers with $t < t^*(p_0)$. Therefore,

$$t^*(p_0)q_0 - p_0 = 0. \tag{10}$$

With regard to the other end of the distribution, we know from Proposition 1 that no buyer will purchase at a price higher than the price chosen by t_1 . Therefore,

$$t^*(p_1) = t_1. \tag{11}$$

For sellers there would appear to be two possibilities. Either the quality of car at p_1 is less than q_1 , in which case sellers of all higher quality cars must have an even higher reservation price, or else the quality at p_1 equals q_1 , in which case p_1 must exceed the reservation price tq_1 . Because all sellers have the same utility function, however, the first possibility can be eliminated. First note that since $t^*(p)$ is increasing, (9) and (10) imply that $t^*(p) > t$ for all p . Integrating (7) and using (9) then imply:

$$q^*(p_1) < q^*(p_0) + \frac{1}{t} (p_1 - p_0) < \frac{p_1}{t}.$$

Therefore, the reservation price of the seller at p_1 must be strictly less than p_1 . Consequently,

$$q^*(p_1) = q_1. \tag{12}$$

In the unpublished appendix I demonstrate that (q^*, π^*) is a D.S.E. if and only if it is a solution to equations (6) and (8), subject to constraints (9)–(12). In general, there is a one-parameter family of equilibria with at least one equilibrium corresponding to each initial value of $\pi^*(p_0)$. A typical equilibrium is illustrated in Figure 4. Price is measured on the vertical axis, quality on the right-hand side of the horizontal axis, and the probability of making a sale on the left-hand side.

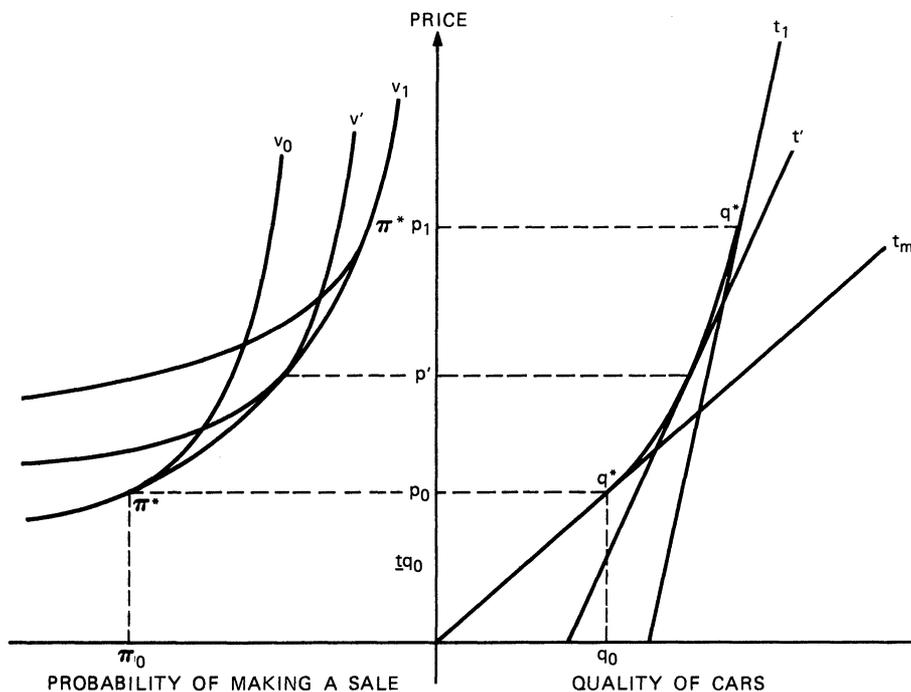
Consider the left-hand side first. The realized probability function, labelled π^* , is a decreasing function of price. At p_0 the indifference curve for sellers of q_0 , labelled v_0 , is just tangent to π^* . For the sellers of q' ($q_0 < q' < q_1$) the indifference curve v' is flatter than v_0 where they intersect and is tangent at $p' > p_0$. The indifference curve for sellers of q_1 is even flatter than v' and is tangent at p_1 . Since π^* is less convex than any of the indifference curves, all points of tangency define the optimal prices for those buyers. The set of announced prices, $P(\pi^e)$, is therefore equal to $[p_0, p_1]$.

On the right-hand side, the realized quality function, labelled q^* , is an increasing function of price. It is defined only on $[p_0, p_1]$. Since $q^*(p_0) = q_0$, the lowest utility index of any active buyer is $t_m = p_0/q_0$. The indifference curve for t_m is just tangent to q^* at p_0 . It also passes through the origin. For buyer t' ($t_m < t' < t_1$), the indifference curve is tangent to q^* at p' , and

¹¹ Note, in particular, that (9) and (12) imply that all sellers must announce a price in equilibrium. In a more general model where sellers are distributed according to their utility index as well as the quality of car they own, this need not be true.

FIGURE 4

A DISCRIMINATING WEAK EXPECTATIONS EQUILIBRIUM



the indifference curve for t_1 is tangent at p_1 . Note that since q^* is convex, these points of tangency also define maxima for the buyers.

The figure, of course, does not show that the resulting distribution of buyers and sellers generates the realized probability function π^* . That depends on the density functions f and h . However, by examining conditions (6)–(12), we can see why we should expect a one-parameter family of equilibria to emerge. Fix any value π_0 between 0 and 1. Now choose a value of p_0 such that $p_0 > tq_0$. Then letting $\pi^*(p_0) = \pi_0$, conditions (9) and (10) determine a unique solution to equations (6)–(8). This need not be an equilibrium. However, by adjusting the value of p_0 , a p_1 can always be found that satisfies (11) and (12). This solution will be an equilibrium. Since π_0 was chosen arbitrarily, there is a distinct equilibrium for each value of π_0 . A formal proof is presented in the unpublished appendix.

□ **Robustness of the sellers' equilibrium.** Let me conclude this section with a brief discussion of the robustness of the sellers' equilibrium. Up to this point, I have been concerned with describing those patterns of expectations that lead to outcomes that do not *explicitly* invalidate those expectations. As we have seen, the set of these expectation patterns can be quite large. Not only can the equilibrium take on a number of different forms, but it can be characterized by almost any level of excess supply. One explanation for the multiplicity of equilibria is the absence of restrictions on the expectations of agents outside the set of prices actually announced. We would like the equilibrium to have the property that even if sellers occasionally experiment by announcing a new

price, their experience will never lead them to revise their expectations in such a way that they permanently alter their equilibrium behavior. In fact, if we look back to the analysis of the buyers' equilibrium, it was precisely this requirement that generated a unique equilibrium. Unfortunately, when we assume that sellers set the price, this approach runs into some new difficulties.

The problem is that it is no longer obvious what it means to restrict expectations to be "correct" at prices at which no trade takes place. Unlike the case where buyers set the price, the quality of cars offered for sale at any price is no longer independent of the sellers' expectations. Whether or not a seller will announce a new price depends upon his expected likelihood of selling a car at that price. But the "correct" probability of selling a car at a new price depends, in turn, on how the buyers revise their expectations about the quality of cars forthcoming at each price. Because of this interaction between the expectations of the agents and their market response on *both* sides of the market, it is no longer obvious what the outcome of experimentation by sellers would be.

Although I believe these problems preclude any clear answer to the question of robustness of the sellers' equilibrium, some of the equilibria may be eliminated with relatively mild restrictions on the expectations of agents. One possible approach is presented in the unpublished appendix. My analysis there suggests no particular reason why experimentation by sellers should bias the market toward either a single price or a discriminating sellers equilibrium. However, I am able to eliminate many of the equilibria with very high levels of excess supply. In fact, for the discriminating equilibria, I argue that robustness requires that excess supply must be zero at the lowest announced price. For single-price equilibria, there will generally be an interval of robust equilibria. This interval need not always include a Walrasian equilibrium, and in some exceptional cases, it may even be empty. It is perhaps in these instances when a distribution of prices is most likely to emerge.

In addition to the question of the robustness of an equilibrium within a given price-setting convention, there is also the question of whether or not the price-setting convention itself is stable. Throughout the paper, I have always assumed that the specific price-setting convention adopted by the market is predetermined, so that each agent is constrained to be either a price taker or a price maker. However, if price takers have an incentive to act as price makers, then it may be reasonable to argue that the price-setting convention will not be followed. I also address this issue in some detail in the unpublished appendix. I argue there that in a buyers' equilibrium no seller can benefit from announcing his own price, but that in a sellers' equilibrium (particularly a D.S.E.), some buyers can generally benefit by announcing their own price and accepting bids. This suggests that the presence of adverse selection may induce a bias towards the market's adopting a convention in which the buyers act as price setters.

6. An overview and some welfare comparisons

■ Having analyzed each of the three different price-setting conventions in some detail, I now present a brief overview of the central results. In the course of the discussion, I shall also indicate some of the welfare conclusions which follow from the analysis.

The basic model specifies a distribution of potential buyers according to their preferences for car quality and a distribution of potential sellers according to the quality of cars they own. For simplicity, the reservation price of each seller is assumed to be strictly proportional to the quality of his car. Using this model, I addressed the following question: Assuming that buyers are able to observe only the average quality of the cars at each price, how does the convention by which prices are set affect the equilibrium allocation of cars and the prices at which they are sold?

The first convention I examined was the standard paradigm of a Walrasian auctioneer. Both buyers and sellers act as price takers and the auctioneer sets a single price that equates supply and demand. Two significant results emerged from the analysis of this problem. First, since the average quality of cars offered for sale may increase with the price, it is possible that higher prices may attract more rather than fewer buyers, thereby generating an upward sloping demand function. As a consequence, there may be multiple equilibria in some cases. Second, when there is more than one equilibrium price, the highest equilibrium price is always Pareto preferred. For the sellers the result is obvious, since there is no rationing in a Walrasian equilibrium. For the buyers it is a consequence of the assumption, implicit in the specification of the utility function, that buyers with higher reservation prices also have a higher marginal rate of substitution of quality for price. This implies that if average quality increases sufficiently with an increase in price to make the marginal buyer better off, it must make all inframarginal buyers better off as well.

The second convention I examined assumed that each buyer explicitly sets his own price. Sellers are then free to submit bids to as many buyers as they wish and to sell at the highest price at which their bids are accepted. Any excess supply is rationed at random. The key to analyzing the equilibrium under this convention is to recognize that because sellers may costlessly search for the highest price, the average quality of cars offered for sale at any price is the same as when an auctioneer sets the price. In both cases it is equal to the average quality of cars owned by those sellers with a reservation value less than the price. The distinction between these two conventions, therefore, lies only in the criteria by which the price or prices are set.

Under a Walrasian equilibrium, a single price is adjusted to equate supply and demand. No account is taken of the possibility that some or all buyers may prefer to purchase at a higher price to obtain higher quality cars. When buyers set the price, this consideration moves to the forefront. Although buyers are constrained to announce prices high enough to attract sellers into the market and away from other buyers, the prices they announce need not necessarily equate demand and supply. All buyers will announce a Walrasian equilibrium price if and only if no buyer would prefer to buy at a higher price. In any other case the equilibrium will be characterized by a distribution of prices with excess supply at all but the lowest price.

It should be noted that whenever the buyers' equilibrium is characterized by price dispersion, no buyer is made worse off relative to the Walrasian equilibrium, since the higher prices attract additional sellers and this, in turn, generates excess supply at the W.E. price. This requires the equilibrium cutoff price to fall below the W.E. price. Since the average quality function remains unchanged and since each buyer still has the option of purchasing at the W.E. price, it then follows that buyers can only benefit by announcing their buyers' equilibrium prices.

For the sellers, the welfare comparison between a Walrasian equilibrium and a buyers' equilibrium is less clearcut. Sellers who enter the market only at prices above the Walrasian price obviously prefer the buyers' equilibrium. Because the effect of attracting these sellers into the market is to create an excess supply of cars at the W.E. price, however, some sellers will be forced to sell at a lower price. Consequently, depending on the distribution of buyers and sellers, the expected utility of some sellers may be lower under the buyers' equilibrium than under the Walrasian equilibrium.

The third convention I examined assumed that sellers set the price. The fundamental distinction between this convention and the previous one is that now each seller must commit himself to a single price which he cannot change, even in the face of excess supply. To the extent that sellers perceive the possibility of an excess supply of cars at some prices, this considerably complicates their decision problem. They must first form an expectation about the probability of selling at each price and then choose the price that maximizes their expected utility. The buyers' problem also becomes more complicated. Since sellers must commit themselves to a single price, the quality of cars at each price need have no relation to the Walrasian average quality function. Buyers must, therefore, form their own expectation about the quality of cars supplied at each price, and then bid at that price which maximizes their expected utility. The market is in equilibrium when the expectations of both the buyers and sellers are confirmed by the market.

This interaction between the expectations of agents and their market response leads to a multiplicity of forms an equilibrium can take. Two of these were discussed in the text. If sellers are sufficiently pessimistic about the likelihood of attracting sellers at a higher price, the only possible equilibrium is one in which all trade takes place at a single price. This may be a Walrasian equilibrium price, but it may also be any other price at which there is an excess supply of cars. In any case, the welfare analysis is exactly the same as if an auctioneer were to announce that price. To the extent that higher prices raise the average quality of the cars, some or all of the buyers may prefer the higher price equilibria. The gain to a seller depends upon how the probability of selling changes and the value of his reservation price. If the ratio of demand to supply is relatively price inelastic, it is possible that every seller would prefer to face a higher price with excess supply than a lower price where the market clears.

I argued that there is also a one-parameter family of equilibria in which sellers of different quality cars each announce distinct prices. This is possible because only sellers with higher quality cars are willing to accept a lower probability for selling to receive a higher price. Consequently, it is possible to adjust the expected probability function so that sellers of higher quality cars choose to announce higher prices, thereby generating an upward sloping average quality function. An equilibrium is then attained when the expected probability function is adjusted so that sellers announce prices that induce the different buyers to bid at different prices in such a way that the expectations of the sellers are realized.

There is apparently very little one can say in the way of welfare comparisons between the discriminating sellers' equilibria and the single price equilibria or the buyers' equilibria. Depending on the distribution of buyers and sellers, there may be no relationship between the equilibrium distribution of prices in the two cases or the average quality of cars supplied at each price. Even among the different discriminating equilibria, strong welfare comparisons do not appear

possible. About all that can be said is that under a discriminating sellers' equilibrium, each owner must sell with some probability at some price. To the extent that owners of the highest quality cars do not enter the market under the other price-setting conventions, therefore, those sellers may benefit from announcing their own prices. But even this result depends on the assumptions that all owners have an identical utility function.

All of these equilibria are, of course, distinct from a full-information competitive equilibrium. Under full information, each quality of car will have a distinct price, with the prices adjusted to equate supply and demand. In equilibrium, cars will be allocated so that the highest quality cars are allocated to those agents with the highest marginal rates of substitution of quality for price. In this respect, the full-information competitive equilibrium is very similar to the discriminating sellers' equilibrium. They differ only because, with asymmetric information, the probability of selling at each price must be adjusted so that owners of lower quality cars *choose* to sell at the lower prices. In general, the assumption of imperfect information on the part of buyers tends unambiguously to increase the welfare of the owners of the lowest quality cars. For all other agents the welfare effects are ambiguous. (Of course, the introduction of imperfect information cannot simultaneously improve the welfare of all agents.)

Let me conclude with a few comments on the scope of these results. I have used the example of the used car market to present these results, because it is a flexible paradigm. It is easy to imagine how different price-setting conventions might operate in this market. However, it is only an example. These results are applicable to many markets where adverse selection is a problem. Of course, not all of the price-setting conventions may be relevant for all markets. When adverse selection is present in a labor market, for instance, the relevant price-setting convention may be to assume buyers announce the price. In contrast for markets where professional or skilled services are for sale (e.g., lawyers, doctors, plumbers), it probably is most accurate to assume that sellers announce the price. If for any reason those sellers with higher quality products have lower marginal rates of substitution between the price they charge and the number of sales they make, the analysis in this paper suggests that the sellers may announce a distribution of prices in which the quality anticipated by buyers at each price is confirmed by the market.

References

- AKERLOF, G. "The Market for 'Lemons': Qualitative Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*, Vol. 84 (August 1970), pp. 488-500.
- STIGLITZ, J. "Prices and Queues at Screening Devices in Competitive Markets." Stanford University IMSS 212, August 1976.
- WILSON, C. "The Nature of Equilibrium in Markets with Adverse Selection." University of Wisconsin, SSRI Discussion Paper 7715, November 1977.
- . "The Welfare Benefits of Price Maintenance in Markets with Adverse Selection." University of Wisconsin, SSRI Discussion Paper 7817, August 1978.