Price and Advertising Signals of Product Quality

Paul Milgrom

Yale University

John Roberts

Stanford University

We present a signaling model, based on ideas of Phillip Nelson, in which both the introductory price and the level of directly "uninformative" advertising or other dissipative marketing expenditures are choice variables and may be used as signals for the initially unobservable quality of a newly introduced experience good. Repeat purchases play a crucial role in our model. A second focus of the paper is on illustrating an approach to refining the set of equilibria in signaling games with multiple potential signals.

Although we economists have included advertising and other selling expenses in various of our models at least since the 1930s, it is only within the last decade or so that we have begun to offer explanations of why advertising might affect customers' choices and thus of why firms might choose to advertise.

The most successfully developed of these models involve firms' using advertising to inform potential customers about the existence, characteristics, and prices of the commodities they offer. This work

We would like to thank Bob Wilson, an anonymous referee, and the audiences at the numerous seminars and conferences at which this work was presented for their interest, comments, and suggestions; CORE for its hospitality; and the National Science Foundation (grants SES 83-08723 to Stanford University and IST 84-11595 to Yale University) and the Institute for Advanced Studies at the Hebrew University of Jerusalem for financial support.

[*Journal of Political Economy,* 1986, vol. 94, no. 4]

has obvious relevance to the huge volume of advertising that is directly informative on these dimensions. Most newspaper advertisements (including especially want ads) would seem to be of this sort.

However, a nontrivial amount of advertising (especially on television) has little or no obvious informational content. A relatively recent example is the ad that was shown when Diet Coca-Cola was introduced: a large concert hall full of people, a long chorus line kicking, a remarkable number of (high-priced) celebrities over whom the camera pans, and a simple announcement that Diet Coke is the reason for this assemblage. Another example from the same period is the advertising campaign for the 1984 Ford Ranger truck, which featured these trucks being thrown out of airplanes (followed by a half dozen sky divers) or driven off high cliffs. These ads carry little or no direct information other than that the product in question exists. But if that is the message being sent, these ads seem an inordinately expensive way to transmit the information. Indeed, the clearest message they carry is, "We are spending an astronomical amount of money on this ad campaign."

In a series of provocative articles, Nelson (1970, 1974, 1978) has suggested that the latter is, in fact, the primary message of such ads and, moreover, that this is a useful, positive message to prospective customers. Nelson differentiated between products on a "search good" versus "experience good" basis. With the former, the relevant characteristics of the product are evident on inspection, and, because there is little gain to misrepresentation, ads for them can be directly informative. With the latter, crucial aspects of the product's quality are impossible to verify except through use of the product. Thus, unless the product is given away, one must buy without really knowing what one is getting. In such a circumstance, a seller's claims to be offering high quality are unverifiable before purchase. In the absence of strong and sure penalties for misrepresentation, such claims can be freely copied. They are consequently meaningless, and consumers will rationally ignore them. As a result, ads for such a product cannot credibly convey much direct information about the product. Yet it remains in the interests of consumers to identify high-quality goods and of the producers of these "best buys" to make themselves known.

Nelson's crucial insight was that the mere fact that a particular brand of an experience good was advertised could be a signal to customers that the brand was of high quality. It is clear that if high-quality brands advertise more and if advertising expenditures are observable (even if not perfectly so), then rational, informed consumers will respond positively to advertising, even if the ads cannot and do not have much direct informational content. What then is needed to complete the explanation is a reason why advertising should be differentially advantageous for high-quality sellers so that they will be willing to advertise at levels that low-quality sellers will not mimic.

The factor on which Nelson focused to provide this linkage was repeat purchases. He argued that, because a high-quality product is more likely to attract repeat purchases, an initial sale is, ceteris paribus, more valuable to a high-quality producer, and such a firm would be willing to expend more—on advertising or whatever—to attract an initial sale. This relationship would then provide the basis for the correlation of quality with the net benefits of signaling that is needed in the standard Spence-type analysis to obtain a separating equilibrium.

Nelson's approach is very insightful and appealing, but it is not worked out in terms of a formal model. Moreover, further consideration reveals what proves to be a major gap in his analysis. Specifically, Nelson did not explicitly treat the pricing decision and the determination of the resultant markup. Yet these are crucial questions.

On the pricing side, if the firm is able to select the price it will charge (subject to whatever competitive pressures may exist), might it not prefer to stimulate sales through its pricing rather than via uninformative ads? Or, if advertising does convince customers of a product's high quality, might not the firm want to alter its price in response to the increased demand? But note that, if such possibilities result in prices that vary systematically with quality, then Nelson's explanation of advertising is undercut. Customers can now infer quality from price and so have no need to look to advertising for a hint as to what quality might be. In this circumstance, why should firms waste money on ads?

But even if the role of prices as possible signals is put aside, their determination remains crucial. This is because the value of an initial sale depends not just on the volume of resulting repeat sales but also on the markups received. Nelson's ceteris paribus assumption apparently means that markups are taken to be the same on high- and low-quality products, and some version of this assumption is clearly indispensable. For if markups were sufficiently greater on low-quality goods (as they would be if prices were the same but production costs were steeply increasing in quality), then the value of an initial sale would be negatively correlated with quality. If customers then responded positively to ads, it would be the low-quality firms that would do the advertising, while if customers understood the incentives facing firms, neither type of firm would advertise.

It thus becomes important to attempt to formulate Nelson's basic ideas in a complete, formal model incorporating both the pricing and advertising decisions. In fact, a number of authors since Nelson have investigated the relationship between quality and the use of the noninformative or image advertising on which he focused, and some have been explicitly interested in formalizing his ideas. However, the issues of both prices and quantities being choice variables that could be signals and of repeat sales being a key phenomenon have not, to our knowledge, been satisfactorily incorporated into a formal analysis.[1]

In this paper we offer a modeling based on the repeat sales mechanism in which both price and advertising are decision variables that may potentially be used as signals of quality.[2] We show that in equilibrium both may simultaneously be used as signals, with the chosen levels of both prices and advertising differing between high- and low-quality firms (and, moreover, differing for the high-quality firms from the levels that would be chosen in the absence of the informational asymmetry about quality). This means, in particular, that customers could in fact infer product quality from observing either price or advertising volume. However, if a high-quality firm were to cut back on either dimension—price or advertising—of its signaling and move the relevant variable toward its full-information optimal level, then a low-quality firm would be willing to mimic, the signal would no longer be credible, and customers would ignore it. Meanwhile, in such an equilibrium the firm uses both variables to signal (rather than just one) since this achieves the desired differentiation at minimal cost. A corollary of this result is that an effective ban on purely dissipative signals (such as advertising is here) may lead to a Pareto-worsening in the allocation of resources.

Three points are worth noting here. First, while we will consistently refer to advertising, the analysis clearly applies to any observable expenditure that does not directly provide information or otherwise

[1] Kihlstrom and Riordan (1984) present an interesting model of advertising as a signal. In their model, however, firms do not choose prices. Instead, a firm's advertising alone determines whether customers believe it to be high- or low-quality, and once this assignment to one or the other submarket is made, prices are determined via a standard supply and demand model. In equilibrium, prices in fact end up being correlated with quality but are not used to infer quality. Schmalensee (1978) offers a model in which consumers follow a rule of thumb. In it, low-quality producers may do the advertising because markups are negatively correlated with quality and customers do not recognize the negative advertising-quality relationship. Johnsen (1976) directly attempts to formalize Nelson's argument but does not obtain existence of equilibrium when both prices and ad budgets are choice variables. (We are grateful to Ed Prescott for this reference.)

[2] Klein and Leffler (1981) offer an alternative, complementary explanation for introductory advertising. In their formulation, unlike ours, quality is a choice variable, and the problem is to motivate firms not to cheat by cutting quality. The incentive to maintain quality comes through positive markups and repeat sales, which are lost once cheating is discovered. However, these profits must be reconciled with free entry. This is achieved by requiring new firms to sink resources on ads in an amount equal to expected operating profits before they can attract any business (see also Shapiro 1983).

improve demand or costs. A shop in a high-rent location or highly visible corporate social responsibility activities are obvious examples. Second, the analysis is strictly applicable only to new products whose quality is not generally known. Thus it says little about advertising for established brands. Third, we emphasize that quality in this analysis is not treated as a choice variable but rather as exogenously given. The problem is not the moral hazard one that the firm may have incentives to cheat by cutting quality. Indeed, we do not even assume that lower quality is necessarily cheaper to produce. It is thus probably best to think of our model as one in which the firm's R & D effort has generated a product of some particular given quality that the firm must decide how to introduce.

While the primary purpose of this paper is to study the role of pricing and advertising for newly introduced experience goods, it may also offer some methodological contribution through providing the analysis of multiple variables being used simultaneously to signal for a single unobservable variable and through illustrating a method of obtaining a "small" set of equilibria and even uniqueness in signaling situations modeled as games.

The second of these actually underlies the first. Models based on games of incomplete information, and signaling models in particular, have typically suffered from an embarrassing plethora of (Nash) equilibria. Not only are there often both pooling and separating equilibria (as well as partial-pooling ones), but also there are typically a horde of each of these types. The source of this multiplicity is the indeterminacy of the inferences that individuals draw "off the equilibrium path," that is, when they see a level of the signal that they can tell ought not to have arisen in equilibrium. Bayes's rule gives no guidance in such situations, and the usual equilibrium notions are unspecific about how such inferences should be made. Yet the beliefs that are formed off the putative equilibrium path and the actions that they generate determine what individuals can accomplish by deviating from the prescribed strategies. They are thus crucial determinants of what will, in fact, be equilibrium behavior. The assumption of equilibrium thus places relatively few restrictions on behavior, and, consequently, many different behavior patterns can be supported in equilibrium.

A long and growing list of authors have addressed the problem of paring down the set of Nash equilibria in signaling models by restricting the allowable beliefs.[3] The approach we use here is first to work

[3] Among the relevant references are papers in which the arguments rely on economic intuition related to the specific context of signaling (e.g., Riley 1975; Milgrom and Roberts 1982; Engers and Schwartz 1984), others involving systematic, game-

only with sequential equilibria, a refinement of the Nash concept developed by Kreps and Wilson (1982). This both forces us to be explicit about the out-of-equilibrium beliefs and restricts them somewhat. More significantly, we also require that the equilibria be immune to sequential elimination of dominated strategies (Moulin 1979; Pearce 1982)[4] and that they meet a further "intuitive criterion" proposed by Kreps (1984).[5] Both of these conditions serve to restrict beliefs further and in economically reasonable ways. Moreover, Kreps (1984) has shown that they are implied by the concept of strategic stability proposed by Kohlberg and Mertens (1984) for general normal form games.[6]

The use of sequential equilibrium eliminates separating Nash equilibria in which low-quality firms, even though revealed as low-quality, deviate from their full-information optimal price or spend money on ads because customers would otherwise buy even less than the full-information amount. Sequential elimination of dominated strategies then requires that the set of price-advertising pairs taken as indicating high quality be as large as possible, in that if it were any larger it would include choices that would be mimicked by a low-quality firm. It thereby rules out separating equilibria with excessive, inefficient amounts of signaling by the high-quality firm since such sequential equilibria are supported only by the belief that a low-quality producer played a dominated strategy. Finally, application of the Kreps (1984) criterion eliminates any equilibria involving pooling on a price-advertising pair at which an appropriate generalization of Spence's signaling condition is met.[7] Thus the only candidates for

---

theoretic approaches of more general applicability (e.g., Selten 1975; Kreps and Wilson 1982; Kohlberg and Mertens 1984; Banks and Sobel 1985; Cho 1985; Cho and Kreps 1985), and a few in which both lines are pursued (e.g., Kreps 1984).

[4] A strategy is dominated if there is another strategy that yields the player at least as high payoffs against every specification of strategies for the other players and strictly more against some strategies. Note that eliminating the dominated strategies for one player may result in previously undominated strategies for another now being dominated.

[5] The essence of this criterion is that at equilibrium there ought not to exist actions that are not being taken but that, if believed to signal high quality, would be advantageous for the high-quality firm to take but not for the low-quality firm. Such acts should be interpreted as, in fact, signaling high quality, and their existence would then upset the equilibrium.

[6] Kreps (1984) uses the same methods as we do to obtain uniqueness in a Spence-type univariate model of job market signaling. The relationship between the economic and game-theoretic arguments for eliminating various outcomes is made very clear in this highly recommended paper.

[7] This condition involves certain strict inequalities on derivatives of the profit functions. In Sec. II we develop an example in which these conditions hold only almost everywhere. As a result, in the example there is generically a unique separating equilibrium, but there may also be equilibria with pooling on the small set of points at which the conditions fail.

equilibrium that meet all our requirements involve the low-quality firm's picking its full-information optimum and the high-quality firm's doing just enough signaling to distinguish itself.

In this context, the equilibrium choices of prices and advertising are given by the solution of a constrained optimization problem for the high-quality producer. The same conditions on profit functions can be used to show that the solution to the problem is unique and that, if the full-information prices of the two types of producers are not too dissimilar, then the solution involves positive levels of advertising as well as price signaling. Of course, this use of both signals is natural given that they are the choice variables in an optimization problem and that the conditions on the profit functions ensure an interior solution.[8]

I. Price and Advertising Signals: A Diagrammatic Exposition

Most of the key ideas underlying our analysis of multidimensional price and advertising signals for quality can be developed graphically. We do so here under assumptions on the existence and shapes of the function relating the firm's equilibrium profits to its initial price and advertising choices and to its actual and perceived quality. In the next section we will present a detailed analysis of a fully specified model.

Consider a firm that has just developed a new product of which it is the sole producer.[9] The product may be of either high quality (H) or low quality (L). The firm knows the actual, realized quality, but the potential customers do not, and there is no credible direct way by which the firm can provide this information before customers make their initial purchase decisions. The firm's immediate decision variables are the price, P, at which it will introduce the product and the amount, A, that it will spend on introductory advertising over and above whatever level is optimally used to inform potential customers of the good's existence, its price, and its verifiable characteristics. These two variables are shown on the axes of figures 1–6 below.

Customers, after observing P and A, make their initial purchase decisions and, through direct use of the product or communication with users, then gain information about product quality. The firm

[8] See Johnson (1976), Grossman (1981), Hughes (1983), Kohlleppel (1983*a*, 1983*b*), Quinzii and Rochet (1984), Holmström and Weiss (1985), and Wilson (1985) for other models explicitly using multiple signals. The Wilson paper is of particular interest here since it extends our analysis to a continuum of qualities and any finite number of signaling variables.

[9] The assumption of monopoly seems natural in this context, at least in comparison with the perfectly competitive alternative. Treating the intermediate case of oligopoly would involve significant additional problems.

then sets its second-period, postintroductory price and carries out any additional advertising it wishes. Customers observe these choices and make their decisions as to whether to buy again in light of the current prices and the information now available. If there are additional postintroductory periods, this latter pattern is repeated in each.

A full specification of this sequence of possible actions, of the information available at each point, and of the resultant payoffs would yield a game of incomplete information in extensive form. A sequential equilibrium would then involve strategies for each player (firm or customer) giving the choice to be made at each decision point as a function of the information then available,[10] as well as beliefs for the customers at each point about the firm's true quality. The beliefs would have to be consistent with the information structure of the problem and, to the extent possible, with the hypothesis that the given strategies were being played, and starting from any decision point the strategies would have to be best responses to one another, given the beliefs. (For full details, see Kreps and Wilson [1982].)

In fact, our interest focuses not on the whole play of the game, as given by the full equilibrium strategies, but rather only on the initial equilibrium choices of P and A by the firm and on the resulting customer beliefs. To study these, it will be enough to assume that each such choice (P, A) induces a unique equilibrium pattern of customer beliefs—represented by the probability $\rho(P, A)$ that is assigned to the firm's producing high quality—and that together these induce a uniquely defined expected present value for the firm's profits throughout the game.

Let $\Pi(P, q, \rho) - A$ denote the function giving the expected present value of the profits to a firm of true quality q ($q = L$ or H) that sets an introductory price of P, spends A on introductory advertising, and is believed with probability $\rho = \rho(P, A)$ to be producing quality H. Note that advertising here has no direct impact on demand or gross profits. Its only possible influence is through prepurchase perceptions of quality. It is thus a purely dissipative signal.

In the present context, it is natural to think that initial sales will be increasing in the perceived quality, as modeled by ρ, and that repeat sales will increase in actual quality, q. These conditions hold in the example in the next section, but it may be possible to concoct examples in which they do not hold but the following lines of analysis would apply. In any case, we do not assume that Π increases in actual

[10] For the firm, this information includes its actual quality, so a strategy for the firm may specify different actions depending on what its actual quality is. Given this, it will be convenient to use terminology that might suggest that both high- and low-quality producers actually exist, even though there is only one firm and its actual quality is either definitely H or definitely L.

quality because costs might also depend on q, but we will assume that profits are increasing in ρ.

Of special interest will be situations in which ρ is zero or one, that is, where customers believe they know the true quality. In such cases, it will be convenient to define $\pi(P, q, L) = \Pi(P, q, 0)$ and $\pi(P, q, H) = \Pi(P, q, 1)$. Thus $\pi(P, q, Q)$ denotes the gross profits of a firm of actual quality ("type") q that is initially perceived to be of type Q and sets price P.

If actual quality were known by potential customers before purchase, then $\pi(P, q, q) - A$ would be the relevant profit function net of advertising expenditure for a firm known to be producing quality q. Clearly, the optimal advertising budget in these circumstances is $A = 0$. Denote the optimal value of P for a firm known to be of type q as P_q^q, that is, P_q^q is P_H^H or P_L^L. We call these the "full-information prices."

Under the actual information conditions that initially obtain with experience goods, q and Q may differ. In this context, define P_Q^q as the maximizer of $\pi(P, q, Q)$. We can now give a first answer to the question whether there exists a separating sequential equilibrium of the signaling game, that is, a sequential equilibrium at which the customers can distinguish high- and low-quality firms by the different price-advertising choices they make.

PROPOSITION 1. There exists a separating sequential equilibrium if and only if for some $(P, A) \geq 0$

$$\pi(P, H, H) - \pi(P_L^H, H, L) \geq A \geq \pi(P, L, H) - \pi(P_L^L, L, L). \quad (1)$$

At any separating sequential equilibrium, the high-quality firm chooses a (P, A) satisfying (1), the low-quality firm chooses $(P_L^L, 0)$, and customers' beliefs are given by $\rho(P, A) = 1$ for the point chosen by the high-quality firm, $\rho(P_L^L, 0) = 0$, and, for all other (P', A'), $\rho(P', A')$ sufficiently small (e.g., zero) that neither player wishes to deviate to (P', A').

The inequalities (1) assert that a high-quality firm would rather choose (P, A) and be perceived as high-quality than be perceived as low-quality and optimize accordingly, while the low-quality firm has the reverse preference. Whatever choice (P_L, A_L) is made by the L, in separating sequential equilibrium this choice must yield $\rho(P_L, A_L) = 0$, and the best such choice is $(P_L^L, 0)$.

The situation in which a separating equilibrium exists is depicted in figure 1. Note that points *under* a given isoprofit curve correspond to *higher* levels of profit. Thus the inequalities hold over the indicated region, and each point in the region corresponds to at least one separating sequential equilibrium.

From this we see that there are typically many separating sequential

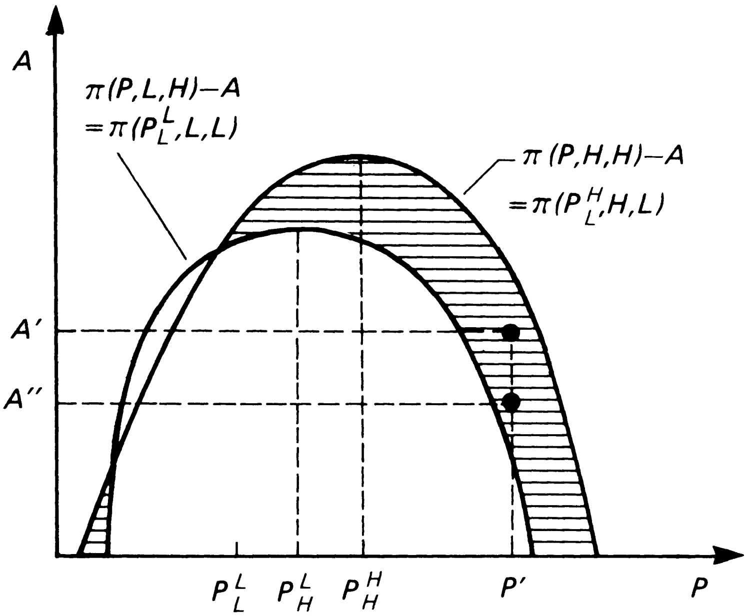

FIG. 1.—Existence of separating equilibrium

equilibria. However, most of these equilibria involve customer beliefs that are arguably implausible. For example, the point (P', A') corresponds to an equilibrium only because customers believe that a firm choosing (P', A'') is likely to be a low-quality producer, even though such a choice is dominated for an L. The best that can happen if an L chooses (P', A'') is that it is taken for an H, but this is worse than the worst that can happen when $(P_L^L, 0)$ is chosen, namely, that the firm is taken for an L. If customers believe that firms do not make dominated choices, then $\rho(P', A'')$ must be one, and the equilibrium where H chooses (P', A') is overturned.

More generally, we shall want to limit our attention to equilibria that remain equilibria even after dominated strategies are removed sequentially from the game.[11] Immunity to sequential elimination of dominated strategies means not only that such strategies are never played (though they could be in Nash or sequential equilibrium) but, more significantly, that the "off-the-equilibrium path" beliefs assign zero probability to such strategies whenever possible. In particular, if a (P, A) pair would necessarily represent play of a dominated strategy

[11] The discussion here will involve only simple elimination of dominated strategies for the firm because we are working with a reduced-form profit function. In general, the appropriate notion is sequential elimination of dominated strategies in the full game.

for one type of firm but not for the other, beliefs following observation of such a choice must ascribe zero weight to the type for which the strategy is dominated. This economically natural condition does not follow from sequentiality, let alone from the Nash assumption. Its absence could result in signaling at very high levels only because lesser amounts are interpreted as (perhaps) indicating that the firm is an L, even though adopting such a choice would be dominated for an L.

Henceforth, we shall reserve the term "equilibrium" for sequential equilibria with this immunity property.

PROPOSITION 2. There exists a separating equilibrium if and only if there is some (P, A) such that (1) holds. At any separating equilibrium, the choice (P_H, A_H) of the high-quality firm must be a solution to

$$\max_{P,A} \pi(P, H, H) - A$$
$$\text{subject to } \pi(P, L, H) - A \leq \pi(P_L^L, L, L), \quad P, A \geq 0. \tag{2}$$

If the solution (P^*, A^*) to (2) is such that $A^* > 0$, then P^* solves

$$\max_{P} \pi(P, H, H) - \pi(P, L, H)$$
$$\text{subject to } \pi(P, L, H) - \pi(P_L^L, L, L) > 0. \tag{3}$$

With elimination of dominated strategies, $\rho(P, A) = 1$ at a separating equilibrium precisely for those points lying on or above the curve $A(P)$ defined by the isoprofit curve $\pi(P, L, H) - A = \pi(P_L^L, L, L)$ when this yields $A \geq 0$, and $A(P) = 0$ otherwise. It must be one for such points for the reasons given above. It cannot be one for points below the $A(P)$ curve, for then an L would make such a choice, be taken as an H, and overturn the equilibrium. The choice that an H makes at a separating equilibrium must yield $\rho = 1$, and since it is free to make any such choice, the first part of the proposition follows. Intuitively, the Lagrange multiplier of the constraint in (2) measures the marginal benefit of advertising. When $A^* > 0$, it must equal the marginal cost, which is unity. Using this observation, the second characterization in the proposition can be derived from the first. Finally, the set of points satisfying (1) is closed and so, if nonempty, contains points satisfying the constraints in (2). This gives the existence result. Of course, whenever there is a unique solution to (2), there is only one separating equilibrium in the game.

Let us now examine some specific cases to see when equilibria involving advertising are most likely to arise. For simplicity, we shall henceforth assume that $\pi(P, L, H)$ is strictly concave in P and that $A(P)$ is positive on an interval $(\underline{P}, \overline{P})$ with $\underline{P} > 0$.

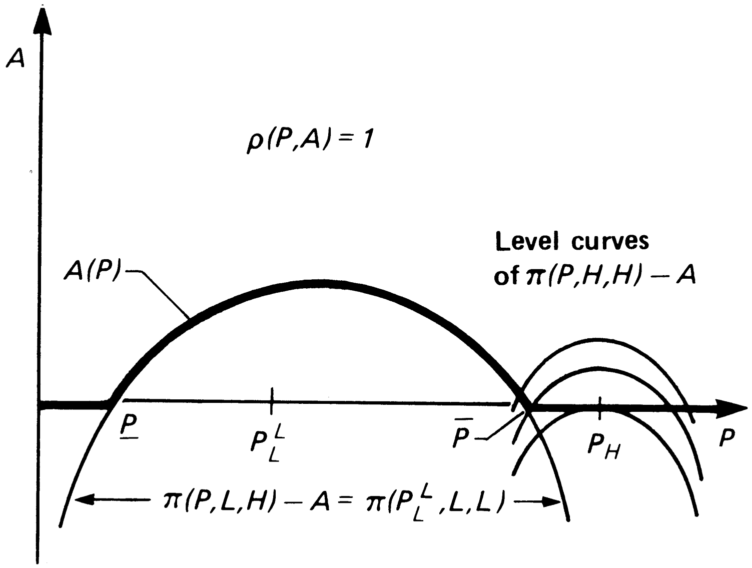

FIG. 2.—Condition (4) fails; equilibrium at the full-information price

Suppose first that $P_H^H \notin (\underline{P}, \bar{P})$; this is the case in which the "natural," full-information price differential is enough to render mimicry unprofitable. Figure 2 illustrates one such case, in which $P_H^H > \bar{P}$. This may occur, for example, when a new high-quality product is very expensive to produce and is aimed at a limited market. A low-quality, low-cost, mass-market producer may be unwilling to limit itself to the small "upscale" market, despite its high margins per unit, so that pricing at P_H^H is enough alone to signal high quality. The case in which $P_H^H < \underline{P}$ has a similar interpretation: When the new high-quality product is very cheap to produce and is aimed at a mass market, the introducing firm may set a low initial price or give away free samples in launching the product. It thereby establishes a large base for repeat sales, which would not be worthwhile if its quality were low. In either case, no costly signaling and, in particular, no advertising are practiced.

Thus a necessary condition for advertising to occur at equilibrium is $P_H^H \in (\underline{P}, \bar{P})$ or, equivalently,

$$\pi(P_H^H, L, H) > \pi(P_L^L, L, L). \tag{4}$$

This condition says that an L would willingly set its price at P_H^H if by so doing it could change its perceived quality from L to H.

Condition (4) leads to signaling but is not by itself sufficient to ensure that the solution to problem (2) involves a positive level of advertising. If advertising is to occur at equilibrium, the solution to problem (2) must occur at a tangency (P^T, A^T) between the isoprofit curves of $\pi(P, L, H) - A$ and $\pi(P, H, H) - A$ with $P^T \in (\underline{P}, \bar{P})$ and $A > 0$. (Note that the isoprofit curves of each type of firm q are vertically parallel with slope equal to $\partial\pi[P, q, H]/\partial P$.) As a necessary condition for an equilibrium with advertising, we thus require that there exist a price P^T at which the isoprofit curves for the two types of firms (both perceived as being high-quality) are tangent and which satisfies

$$\pi(P^T, L, H) > \pi(P_L^L, L, L). \tag{5}$$

Condition (5) is equivalent to requiring that $P^T \in (\underline{P}, \bar{P})$ so that $A(P^T) > 0$.

Note, however, that even if P^T exists the second-order conditions may fail. This leads us to consider the curvature properties of the objective function in (3):

$$\pi(P, H, H) - \pi(P, L, H) \text{ is pseudoconcave in } P; \tag{6}$$

$$\pi(P, H, H) - \pi(P, L, H) \text{ is strictly pseudoconvex in } P. \tag{7}$$

Condition (6) implies that a tangency point is in fact a maximum for problem (3), while (7) means that such a point is actually a minimum.

PROPOSITION 3. Assume that there exists some (P, A) for which (1) holds so that a separating equilibrium exists. If conditions (4)–(6) hold,[12] then there is a separating equilibrium with positive advertising. If the condition in (6) is replaced by strict pseudoconcavity, there is a unique separating equilibrium. If either (4) or (5) fails or if (7) holds, then all separating equilibria have zero advertising.

Think of the H's problem as one of selecting P and A optimally while imposing enough costs on an L making the same choice to render such mimicry unprofitable. If condition (4) holds, then this choice must deviate from the unconstrained optimum at $(P_H^H, 0)$ and so must involve costs for the H. Solving the problem then involves considering the relative effects on each type of firm of changing P and A. Changes in A always affect the profits of either type equally, but the effects of changes in P are more subtle. Condition (6) can be interpreted as saying that movements of P toward P^T involve less cost (or confer more benefits) for an H than for an L, while under (7) movements away from P^T are always cheaper for an H than an L. Suppose now that conditions (5) and (6) hold. Then the costs of price adjustment are equalized for the two types at P^T and further adjustments in P are more costly for the H than for the L, but the necessary differentiation has not been achieved. The H then prefers to distin-

[12] Note that (5) and (6) together actually imply (4).

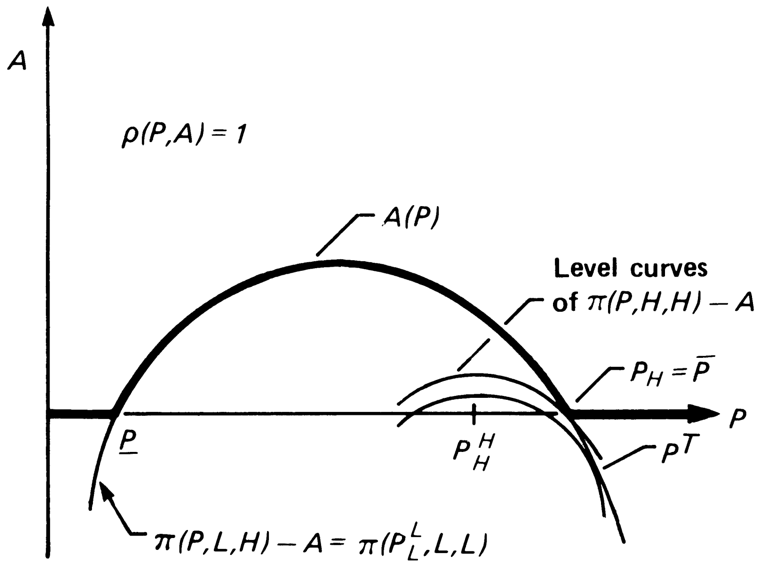

FIG. 3.—Condition (5) fails; equilibrium at ($\bar{P}$, 0)

guish itself by increasing advertising rather than adjusting price further. On the other hand, if (5) fails, then price adjustments (toward P^T) are cheaper for the H than for the L over the entire relevant range and no advertising is employed (fig. 3), while if (7) holds, then the H's problem is again solved at a boundary point of the interval [$\underline{P}$, $\bar{P}$] (see fig. 4).

Figures 5 and 6 illustrate possible cases with positive advertising.

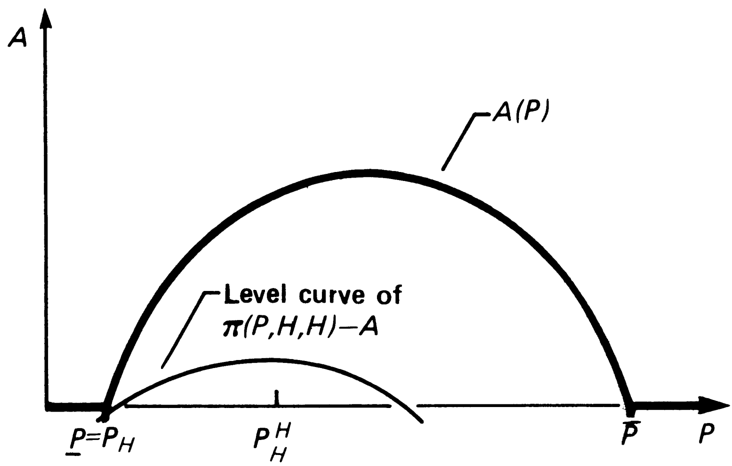

FIG. 4.—Condition (7) holds; equilibrium at ($\underline{P}$, 0)

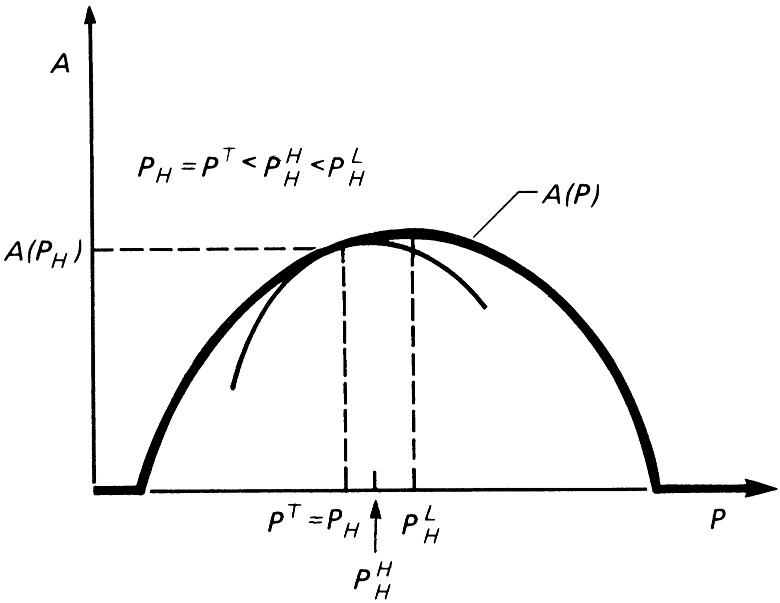

FIG. 5.—Equilibrium with positive advertising and lowered price ($P_H < P_H^H$)

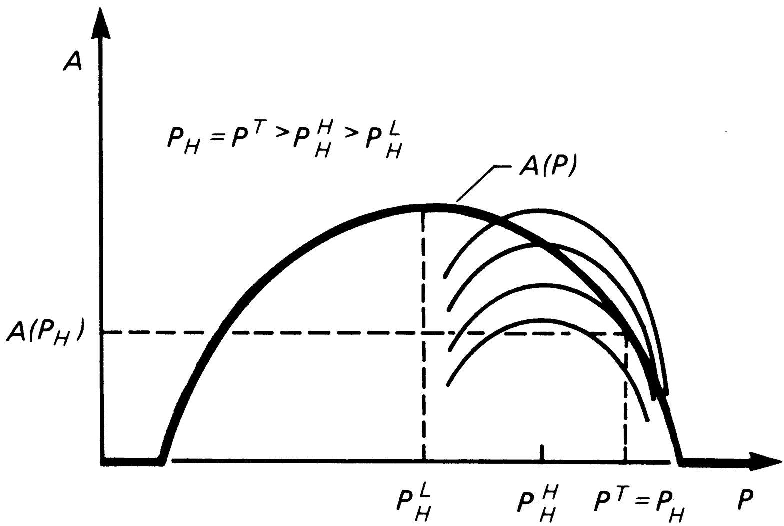

FIG. 6.—Equilibrium with positive advertising and raised price ($P_H > P_H^H$)

We see that price might be raised or lowered from P_H^H for signaling purposes, depending on whether P_H^H exceeds or falls short of P_H^L. If these values agree (and so also equal P^T), then all signaling is via advertising.

In summary, the solution to the optimization problem and, under (1), the choice of the H in separating equilibrium is $(P_H^H, 0)$ if $P_H^H \notin (\underline{P}, \overline{P})$, is $(\underline{P}, 0)$ or $(\overline{P}, 0)$ if P_H^H is "too close" to $\underline{P}$ or $\overline{P}$ (i.e., [5] fails) or if condition (7) holds, and is (P^T, A^T) if (4)–(6) hold. The questions that remain are whether the conditions of (1) are met, so that the separating equilibrium exists, and whether there are other, nonseparating, equilibria.

In fact, there may be pooling equilibria in this game, that is, equilibria in which both the H and the L select the same point (P, A) with positive probability. Indeed, when there is no point (P, A) satisfying (1), separating equilibria will not exist, and any equilibrium will involve pooling. Nevertheless, one may argue on various grounds that pooling equilibria are implausible for this model, at least when separating equilibria exist. For one thing, they make much more severe informational demands on the customers than do separating equilibria. To play their separating equilibrium strategies, customers need know only $A(P)$, which determines the set of things a low-quality producer would be willing to do to be thought high-quality and thus what choices credibly signal high quality. The customers need not know the profit function of the H, nor need they agree on the probability that a given producer is an H, as they must to play their parts in a pooling equilibrium.

A recent game-theoretic development, due to Kreps (1984), provides an additional way to criticize the pooling equilibrium, though it does rely on the customers' being very well informed. The idea is this: Fix some pooling equilibrium whose "stability" is to be checked. Suppose that the high-quality firm would strictly prefer to spurn its equilibrium choice and choose some point (P, A) that is not currently being chosen if by so doing it would be thought high-quality. Suppose, too, that a low-quality firm would rather adhere to the equilibrium price and advertising levels and be thought to be high-quality with the given equilibrium probability than play (P, A), no matter what inferences the customers might draw from the observation of (P, A). Then, according to the Kreps criterion, if (P, A) were played, the customers should infer that it is a move by a high-quality firm. The reasoning is that a low-quality firm could never benefit from such a move and thus would not even experiment with it, but a high-quality firm just might. However, if customers formed their beliefs this way, then an H could profitably separate itself, and the pooling equilibrium would not survive.

Kreps has applied this criterion to a game-theoretic version of Spence's (1973) labor market signaling model. He found that the criterion rules out all pooling equilibria. This finding relies only on the now-familiar single-crossing property of indifference curves that holds in Spence's model. This property reflects the hypothesis that the cost of signaling by acquiring education is inversely related to a worker's ability. The corresponding condition in our model would be that an H always be more willing to signal than an L in return for a given change in its perceived quality. What signaling means in our multidimensional, nonmonotonic context is, however, complicated. Intuitively, the appropriate condition is that there exist some feasible direction in price-advertising space such that it takes a smaller increase in perceived quality to compensate a high-quality firm than a low-quality firm for a move in that direction. To express this as a condition on marginal rates of substitution, we assume that $\Pi(P, q, \rho)$ is continuously differentiable with $\partial\Pi/\partial\rho > 0$. Then the appropriate condition is: For any P and ρ, one or more of the following three inequalities holds:

$$\left(\frac{\partial\pi^H}{\partial P}\right)\left(\frac{\partial\pi^H}{\partial P} - \frac{\partial\pi^L}{\partial P}\right) > 0, \tag{8a}$$

or

$$\frac{\partial\pi^H}{\partial\rho} - \frac{\partial\pi^L}{\partial\rho} > 0, \tag{8b}$$

or

$$\left(\frac{\partial\pi^L}{\partial P}\right)\left(\frac{\partial\pi^L/\partial P}{\partial\pi^L/\partial\rho} - \frac{\partial\pi^H/\partial P}{\partial\pi^H/\partial\rho}\right) > 0, \tag{8c}$$

where all derivatives are evaluated at (P, ρ), and $\pi^q(P, \rho) \equiv \Pi(P, q, \rho)$.

These inequalities state that at any point (P, A, ρ) it is possible to find a direction of change that involves increasing ρ and not decreasing A and that makes the H better off and the L worse off. Inequality (8b) states that there is some small increase in ρ and A alone that does this, while (8c) requires that there exist some increase in ρ and small change in P (which will be costly for an L) that has the desired effect. The first factor in (8c) is positive when it is costly for an L to increase its price and negative when it is costly for an L to reduce it. Finally, inequality (8a) requires that there be some small change in P (which will be profitable for an H) and increase in A that makes an H strictly better off and an L strictly worse off. Certainly, an H would be delighted to make such a change to enhance its perceived quality very slightly, while an L would not.

Two important implications follow from condition (8).

PROPOSITION 4. Assume that (8) holds. Then a separating equilibrium exists and satisfies the Kreps criterion. Moreover, no pooling equilibrium satisfies the Kreps criterion.

Proof. It is immediate from the definitions that any separating equilibrium satisfies the Kreps condition. Thus, for the first part of the proposition, it suffices to prove that such an equilibrium exists. One does this by showing that there exists a path $\{(P(t), A(t), \rho(t)); 0 \leq t < 1\}$ beginning at $(P_L^H, 0, 0)$ and ending with $\rho(1) = 1$, where the profits of an H increase along the path and the profits of an L do not. Then, since $\pi(P_L^H, L, L) \leq \pi(P_L^L, L, L)$, the point $(P(1), A(1))$ satisfies inequality (1), and proposition 2 applies. (Note that, for this argument, the inequalities in [8] can be weak.)

A similar construction is used to prove that there cannot be a pooling equilibrium satisfying the Kreps condition. Suppose there were an equilibrium at which both an H and an L play some (P', A') and are thought high-quality with probability ρ'. If there is a path starting at (P', A', ρ') and ending at some point $(P, A, 1)$ along which the profits of an H increase and those of an L decrease (weakly), then (P, A) overturns the equilibrium according to the Kreps condition. Thus we can prove both parts of the proposition simultaneously by proving that there is such a path starting from any (P', A', ρ').

Fix (P', A', ρ') and let $\bar{\rho}$ be the supremum of $\rho(1)$ over all paths $\{(P(t), A(t), \rho(t))\}$. Suppose $\bar{\rho}$ is actually attained by a path ending at the point $(P, A, \bar{\rho})$ and that $\bar{\rho} < 1$. Then, whichever of (8a)–(8c) holds, there exists a direction $\Delta = (\Delta_P, \Delta_A, \Delta_\rho)$ such that $\Delta_\rho > 0$, $\Delta_A \geq 0$, and $\Pi(P + \epsilon\Delta_P, q, \bar{\rho} + \epsilon\Delta_\rho) - (A + \epsilon\Delta_A)$ is increasing in ϵ when $q = H$ and decreasing when $q = L$. Then the path attaining $\bar{\rho}$ can be extended, contradicting the hypothesis that $\bar{\rho}$ is the supremum over all paths. We omit the purely technical argument that the supremum $\bar{\rho}$ is actually attained by some path. Q.E.D.

Thus, under conditions (4), (5), (6), and (8), we have a unique equilibrium satisfying the Kreps criterion. In it, advertising signals quality, even though this advertising carries no direct information and, in effect, corresponds to a public burning of money. It is used, however, because signaling is worthwhile for the high-quality firm and doing so with both price and advertising is cheaper than with price changes alone. Indeed, suppose that advertising of this sort were banned. Then the firm might well continue to signal but to do so through price alone. In this case, price might rise to $\bar{P}$, which represents a Pareto-worsening: profits fall for an H and are unchanged for an L, while customers end up with the same information as before but pay higher prices.

Because the conditions we have stated here are on an endogenous construct, the Π function, it is not immediately obvious to what extent they can be met in a fully specified model in which repeat sales are explicit. We present such a model in the next section. In it, (4)–(6) hold, but the strict inequality in (8) holds only almost everywhere. We also obtain some comparative statics results that suggest how a firm introducing a new, high-quality product would select its price and advertising levels as its costs varied.

II. Analysis of an Explicit Example

We consider the following specialization of the setup described in the previous section. Quality is operationalized as the probability that a randomly selected customer will find the product satisfactory. Potential customers know only that the product is of one of two possible quality levels, L or H, $1 \geq H > L > 0$, and each initially assigns strictly positive prior probability to each possibility.

The set of potential customers corresponds to an interval $[0, R]$ with a uniform distribution[13] and total mass R. We also assume that the probability of the good's being satisfactory to consumer r is independent of r and that customers have no use for more than one unit of the good per period. A customer with valuation r who buys in a particular period at price P receives utility in that period of $r - P$ if the good turns out to be satisfactory and of $-P$ if it does not. A customer who does not buy receives zero. For simplicity, we assume that customers who do not purchase in the first round can never purchase the good. We also assume that an individual can learn whether the good is satisfactory for him or her from only a single purchase. Both of these assumptions can be relaxed.

The costs of production for a firm producing x units of quality q are $C_q x$.[14] Although one might expect that $C_H > C_L$, we do not assume this (see below). Thus, if the firm sets price P and advertising level A in some period and produces and sells x units of quality q, its profit in that period is $(P - C_q)x - A$.

Both the firm and customers maximize the expected present value of their payoffs up to a common (finite or infinite) horizon, T, using a common discount factor $\delta \in (0, 1)$. Let $\Delta = \Sigma_1^T \delta^t$.

To analyze the game that results from this setup, we assume that each decision maker acts in a sequentially rational fashion, following a

[13] The uniform distribution gives rise to a linear demand, which, in turn, facilitates computation. However, our basic results hold as long as the distribution of reservation prices gives rise to a demand curve showing decreasing marginal revenue.

[14] We ignore any fixed costs because their inclusion would have no effect on the solution.

strategy from each point forward that maximizes his or her expected payoff given his or her current information and beliefs and the conjectured behavior of the others. We further require that these beliefs be formed in a manner consistent (whenever possible) with the initial beliefs and the hypothesis that the observed history of play has been generated by the specified strategies, and that the conjectured behavior be consistent with the actual choices. Thus we are employing a refinement of the Nash equilibrium concept in the spirit of the game-theoretic criteria of perfectness (Selten 1975) or sequentiality (Kreps and Wilson 1982). Moreover, we require that the equilibrium be unaffected by sequentially eliminating dominated strategies from the normal form of the game. (We will impose our final restriction, the Kreps criterion, later.)

With a continuum of customers, as long as individual purchases are not observable, it is clearly a dominant strategy for each to act as a price taker in any period. Given this, it can easily be shown that, once we eliminate dominated strategies, the equilibrium demand facing the firm with true quality q at a price p in any postentry period is $q(R - p)$ if p exceeds the valuation r^* of the marginal customer in the first period and $q(R - r^*)$ otherwise, because only a fraction q of the initial $R - r^*$ purchasers will be satisfied with the good and willing to consider repeat purchases. Thus postintroductory advertising has no value and will not be used. Moreover, postintroductory prices will always weakly exceed r^* in any equilibrium, and so the marginal customer in the initial period will obtain no surplus on repeat sales. Consequently, he or she will make an initial purchase at price P only if it is expected to be worthwhile on its own. If ρ is the probability assigned to the true quality being H, the marginal customer is then defined by $\rho H r^* + (1 - \rho) L r^* = P$. Defining the expected quality $Q(\rho) = \rho H + (1 - \rho)L$, we then have $r^* = P/Q(\rho)$. The individual uncertainty about whether the good will prove satisfactory lowers the willingness to pay below the valuation placed on a unit known to be satisfactory. Thus first-period demand at price P with advertising A is $R - [P/Q(\rho)]$, where $\rho = \rho(P, A)$. These relationships are shown in figure 7.

Given these demand relations, the firm that initially charged P and had perceived quality $Q(\rho)$ will charge a constant price in all postentry periods of $p = \max[P/Q(\rho), m(q)]$, where $m(q) = (R + C_q)/2$ would be the simple monopoly price if the good were sure to be satisfactory. Note that the introductory price P is lower than p as long as $Q(\rho) \neq 1$; that is, introductory discounts are used as long as the good is not perceived as surely satisfactory.

The profits $\Pi(P, q, \rho)$ of a firm of type q charging an introductory

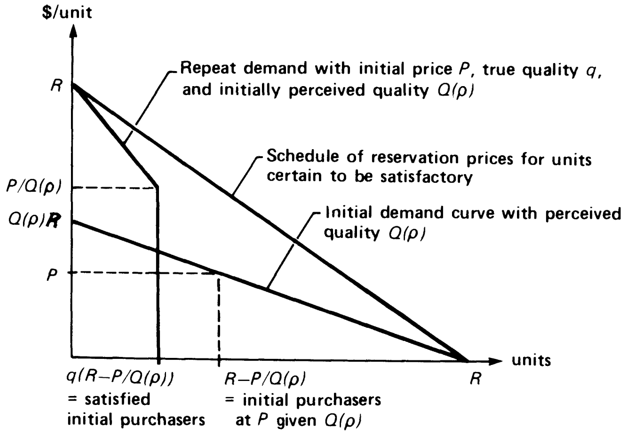

FIG. 7.—Demand relationships in the example

price P and regarded as being a high-quality producer with probability ρ thus are

$$\left[R - \frac{P}{Q(\rho)}\right](P - C_q)$$
$$+ \Delta q\left\{R - \max\left[\frac{P}{Q(\rho)}, m(q)\right]\right\}\left\{\max\left[\frac{P}{Q(\rho)}, m(q)\right] - C_q\right\}.$$

When $\rho = 0$ (respectively 1) $Q(\rho)$ is L (respectively H), and, defining $\pi(P, q, H) = \Pi(P, q, 1)$ and $\pi(P, q, L) = \Pi(P, q, 0)$, we have an exact correspondence with the notation in the previous section.

Using these profit functions, one can routinely check the various conditions assumed in the previous section. First, although (8) does not hold as stated, a variant in which the inequalities are weak is satisfied. As examination of the argument given in support of proposition 4 reveals, this is sufficient to establish the existence of a separating equilibrium in the example.

To examine this equilibrium we must check the other conditions. One easily sees that condition (6) holds if $C_H \geq C_L$. In fact, $\pi(P, H, H) - \pi(P, L, H)$ is strictly concave for $C_H > C_L$, while for $C_H = C_L$ this function is constant for $P < Hm(H)$ and then decreases. When $C_L > C_H$, either (4) fails because $P_H^H < \underline{P}$ or (7) holds because $\partial\pi(P, H, H)/\partial P < \partial\pi(P, L, H)/\partial P$ for all $P \geq \underline{P}$. Moreover, in this latter case we have $\pi(\underline{P}, H, H) > \pi(\overline{P}, H, H)$. For $C_L < C_H < C_L[(H + \Delta H)/(1 + \Delta H)] \times [(1 + \Delta L)/(H + \Delta L)]$ we have $\underline{P} < P^T < P_H^H$, where

$$P^T(\rho) = Q(\rho)\left[\frac{R}{2} + \frac{(1 + \Delta H)C_H - (1 + \Delta L)C_L}{2\Delta(H - L)}\right]$$

is the price at which $\Delta\Pi(P, H, \rho)/\partial P = \partial\Pi(P, L, \rho)/\partial P$ and where $P^T = P^T(1)$.[15] For higher values of C_H, P^T exceeds P_H^H. However, as long as C_H does not exceed C_L by too great an amount, (5) holds and so $P^T < \bar{P}$. Finally, (4) holds for values of C_H lying in an open interval containing C_L.

From this, we see that except when $C_H = C_L$ there is a unique separating equilibrium. For $C_H < C_L$, the equilibrium involves a price $P_H \leq \underline{P}$ and no advertising. Here $P_H = P_H^H < \underline{P}$ if (4) fails, and $P_H = \underline{P} \leq P_H^H$ if (4) holds. At $C_H = C_L$, there is an interval $[\underline{P}, Hm(H)]$ of prices P_H and corresponding advertising levels $A(P_H)$ that constitute equilibrium choices for the H. Note that $A(P_H)$ is positive over this interval except at $P_H = \underline{P}$. For $C_H - C_L$ positive but not too large, the unique solution is at $P_H = P^T$ with $A(P^T) > 0$. In this range, the equilibrium level of advertising first increases, then decreases with the cost and the price. Finally, once C_H is sufficiently large that (5) fails, the solution again involves $A = 0$, but with $P_H \geq \bar{P}$. If (5) fails but (4) still holds, the solution is $P_H = \bar{P}$; if (4) fails as well, the solution is $P_H = P_H^H > \bar{P}$. These relationships are shown in figure 8.

The intuition behind these results is simple. Assume that some signaling activity is necessary, that is, $P_H^H \in (\underline{P}, \bar{P})$. With $C_H < C_L$, the H unambiguously finds price reductions cheaper than does the L: not only is producing to meet the higher initial quantity demanded less expensive for the H, but also a greater fraction of the initial customer base will make repeat purchases that are more profitable for the H as well. Thus the producer of a new, high-quality good that represents a technological breakthrough yielding very low relative costs of production will introduce its product at very low prices, perhaps giving free samples, but will not use advertising as a signal of its quality. When C_H exceeds C_L, the relative effects of price changes on the two types of firm (given that each is perceived as an H) is more complicated. Increasing price and reducing quantity is less costly in terms of profits in the introductory period for the H since it receives a smaller markup at any price. However, the effect on the profits from repeat sales of raising the introductory price might be larger or smaller for the H than for the L. For $P > Hm(H)$, the effect is negative for both, but which is the more negative depends on the relative size of two effects. Raising price costs the H less "per customer" lost, since it receives a lower markup, but the H's larger fraction of repeat purchasing means it loses customers at a greater rate than does the L when P is increased. Unless $P_H^H = P_H^L = P^T$, the overall effects will differ between the two firms at P_H^H, and the H will find a price change to be a cost-

[15] Note that $\partial\Pi(P, H, \rho)/\partial P \gtreqqless \partial\Pi(P, L, \rho)/\partial P$ and $\partial\Pi(P, H, \rho)/\partial\rho \lesseqqgtr \partial\Pi(P, L, \rho)/\partial\rho$ as $P \lesseqqgtr P^T(\rho)$.

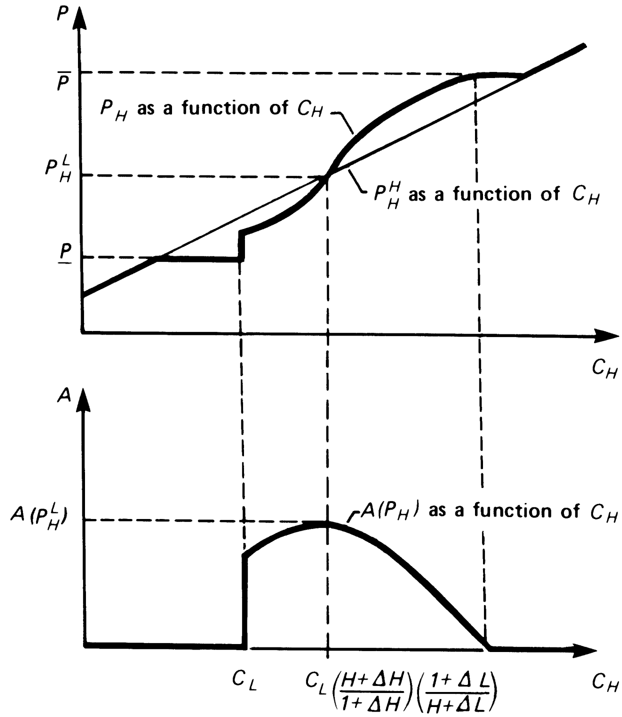

FIG. 8.—Equilibrium values of A and P_H as C_H varies

effective way to signal. Whether advertising is also used then depends on whether the costs of price increases are equalized for the two types before $\bar{P}$ is reached. If they are, advertising results; if not, only price signaling occurs.

While there is generically a unique separating equilibrium in this model (i.e., except when $C_H = C_L$), there will also be a very large set of pooling equilibria. For example, if $\bar{\rho}$ is the common prior assigned to H, then any point (P^*, A^*) with $\Pi(P^*, H, \bar{\rho}) - A^* \geq \max \pi(P, H, H) - A(P)$ and $\Pi(P^*, L, \bar{\rho}) - A^* \geq \pi(P_L^L, L, L)$ can give a pooling equilibrium. These points alone can constitute an open set.

The issue then is whether any of these pooling equilibria survive application of the Kreps criterion. In fact, as the failure of (8) to hold everywhere might suggest, some do. In particular, if there exists a separating equilibrium in which the H has a positive level of advertising, then there also exist pooling equilibria. These are of two forms. First, there may be a family of pure strategy equilibria in which both types choose the same price and advertising level with probability one. The members of this family all involve a common price $P^T(\bar{\rho})$, where $\bar{\rho}$ is the common prior assigned to the firm being an H. The equilibria in this class are indexed by the level of advertising, which must not exceed $A(\bar{\rho}) = \Pi(P^T(\bar{\rho}), H, \bar{\rho}) - [\Pi(P^T, H, 1) - A(P^T)] = \Pi(P^T(\bar{\rho}), L, \bar{\rho}) - \Pi(P_L^L, L, 0)$. These equilibria always include one in which no advertising is used. Their existence depends on $\bar{\rho}$'s being sufficiently high that $A(\bar{\rho}) \geq 0$, that is, on its being sufficiently likely a priori that the firm is high-quality. There will also be a second family of pooling equilibria involving mixed strategies. In these, each type randomizes between its separating equilibrium choice, $(P_L^L, 0)$ or $(P^T, A(P^T))$, and points of the form $(P^T(\rho), A(\rho))$, where $P^T(\cdot)$ and $A(\cdot)$ are as above and ρ is the posterior probability that the firm is an H given that the commonly chosen point is observed. These two classes of pooling equilibria survive application of the Kreps criterion because prices of the form $P^T(\rho)$, $\rho \in [0, 1]$, are precisely and uniquely those at which (8) fails, so that one cannot find a move in (P, A, ρ) space that benefits the H and not the L. Since (8) holds except at prices $P^T(\rho)$, all other pooling equilibria are eliminated.

III. Conclusion

We have constructed a model to formalize Nelson's insight that apparently uninformative advertising for an experience good could be a signal for product quality. In doing so, we have also treated the pricing decision of the firm and allowed for the possibility that price itself might be a signal. Our analysis has confirmed and extended Nelson's fundamental point: advertising may signal quality, but price signaling will also typically occur, and the extent to which each is used depends in a rather complicated way, inter alia, on the difference in costs across qualities.

Interestingly, while our analysis confirms Nelson's fundamental point concerning the role of advertising, its inclusion of the pricing decision upsets the intuition that a high-quality producer will have a higher marginal benefit from attracting an initial sale and that this would provide the basis for the high-quality firm's being willing to advertise more. As noted earlier, all the requisite signaling takes place through the price unless choice of the price P^T does not, by itself, achieve the necessary differentiation. Only in this case is advertising used as a signal. But the present value to a firm perceived to be high-quality of an additional sale achieved through pricing is $(L - H)\partial\pi(P, q, H)/\partial P$, and at P^T this is independent of the firm's actual quality!

The essential difficulty is that the notion of the "marginal benefit to attracting another initial sale" is not well defined: in particular, it depends on who the marginal customer is, that is, on the price being charged and on the beliefs that customers hold. Moreover, once one allows that price might be a choice variable, then ambiguity remains even after the price and beliefs are specified because there are several ways one might generate the sale. If one imagines somehow obtaining an extra sale without changing perceived quality or price, then, at price P^T and perceived quality H, the marginal profit for a firm of type q from an extra sale in the example is $\{P^T[(1 + \Delta q)/H] - C_q(1 + \Delta q)\}$, and this expression is greater for $q = H$ than for $q = L$. If the additional sales are generated through extra advertising, however, then the marginal benefit is greater for the low-quality producer as long as price does not exceed $P^T(\rho)$ and, in particular, if $P \leq P^T = P^T(1)$. If price is the means used to increase sales, then the net marginal benefits from an additional sale are a negative constant, $(L - H)$, times $\partial\pi/\partial P$. If $C_H > C_L$, then, in separating equilibrium with signaling, this marginal benefit is always at least as large for the low-quality producer because the solution always involves $P^H \leq P^T$.

References

Banks, Jeffrey, and Sobel, Joel. "Equilibrium Selection in Signalling Games." Discussion Paper no. 85-9. La Jolla: Univ. California, San Diego, Dept. Econ., 1985.

Cho, In-Koo. "A Refinement of the Sequential Equilibrium Concept." Mimeographed. Stanford, Calif.: Stanford Univ., Grad. School Bus., 1985.

Cho, In-Koo, and Kreps, David M. "More on Signalling Games and Stable Equilibria." Mimeographed. Stanford, Calif.: Stanford Univ., Grad. School Bus., 1985.

Engers, Maxim, and Schwartz, Marius. "Signalling Equilibrium Based on Sensible Beliefs: Limit Pricing under Incomplete Information." Discussion Paper no. 84-4. Washington: Dept. Justice, Antitrust Div., Econ. Policy Office, 1984.

Grossman, Sanford J. "The Informational Role of Warranties and Private Disclosure about Product Quality." *J. Law and Econ.* 24 (December 1981): 461–83.

Holmström, Bengt R., and Weiss, Laurence. "Managerial Incentives, Investment, and Aggregate Implications: Scale Effects." *Rev. Econ. Studies* 52 (July 1985): 403–25.

Hughes, Patricia. "Signaling by Direct Disclosure under Asymmetric Information." Mimeographed. Vancouver: Univ. British Columbia, 1983.

Johnsen, Thore. "Advertising, Market Equilibrium and Information." Ph.D. dissertation, Carnegie-Mellon Univ., 1976.

Kihlstrom, Richard E., and Riordan, Michael H. "Advertising as a Signal." *J.P.E.* 92 (June 1984): 427–50.

Klein, Benjamin, and Leffler, Keith B. "The Role of Market Forces in Assuring Contractual Performance." *J.P.E.* 89 (August 1981): 615–41.

Kohlberg, Elon, and Mertens, Jean-François. "On the Strategic Stability of Equilibria." Working Paper no. 1-785-012. Boston: Harvard Univ., Grad. School Bus. Admin., 1984.

Kohlleppel, Laurenz. "Multidimensional Market Signalling." Discussion Paper no. 125. Bonn: Univ. Bonn, Inst. Gesellschafts- und Wirtschaftswissenschaften, 1983. (*a*)

———. "Properties of Sorting Equilibria." Discussion Paper no. 133. Bonn: Univ. Bonn, Inst. Gesellschafts- und Wirtschaftswissenschaften, 1983. (*b*)

Kreps, David M. "Signalling Games and Stable Equilibria." Research Paper no. 758. Stanford, Calif.: Stanford Univ., Grad. School Bus., 1984.

Kreps, David M., and Wilson, Robert. "Sequential Equilibria." *Econometrica* 50 (July 1982): 863–94.

Milgrom, Paul, and Roberts, John. "Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis." *Econometrica* 50 (March 1982): 443–59.

Moulin, Hervé. "Dominance Solvable Voting Schemes." *Econometrica* 47 (November 1979): 1137–51.

Nelson, Phillip. "Information and Consumer Behavior." *J.P.E.* 78 (March/April 1970): 311–29.

———. "Advertising as Information." *J.P.E.* 81 (July/August 1974): 729–54.

———. "Advertising as Information Once More." In *Issues in Advertising: The Economics of Persuasion,* edited by David G. Tuerck. Washington: American Enterprise Inst., 1978.

Pearce, David. "Ex Ante Equilibrium: Strategic Behavior and the Problem of Perfection." Research Memorandum no. 301. Princeton, N.J.: Princeton Univ., Dept. Econ., 1982.

Quinzii, Martine, and Rochet, Jean-Charles. "Multidimensional Signaling." Technical Report no. 453. Stanford, Calif.: Stanford Univ., Inst. Math. Studies Soc. Sci., 1984.

Riley, John G. "Competitive Signalling." *J. Econ. Theory* 10 (April 1975): 174–86.

Schmalensee, Richard. "A Model of Advertising and Product Quality." *J.P.E.* 86 (June 1978): 485–503.

Selten, Reinhard. "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games." *Internat. J. Game Theory* 4, no. 1 (1975): 25–55.

Shapiro, Carl. "Premiums for High Quality Products as Returns to Reputations." *Q.J.E.* 98 (November 1983): 659–79.

Spence, Michael. "Job Market Signaling." *Q.J.E.* 87 (August 1973): 355–74.

Wilson, Robert. "Multi-dimensional Signalling." *Econ. Letters* 19, no. 1 (1985): 17–21.