Rochester Center for
Economic Research

Recurrent Advertising

Horstman, Ignatius J. and Glenn M. MacDonald

Working Paper No. 105
October 1987

University of
Rochester

Preliminary

RECURRENT ADVERTISING

by

Ignatius J. Horstmann*

and

Glenn M. MacDonald+

Working Paper No. 105

September 1987

An early version of this paper, entitled "Truth in Advertising," was presented at the 1986 Winter Meetings of the Econometric Society, New Orleans.

*Centre for Decision Sciences and Econometrics, U.W.O. and Rochester Center for Economic Research, University of Rochester.

+Centre for Decision Sciences and Econometrics, U.W.O.; Rochester Center for Economic Research, University of Rochester; and Economics Research Center/NORC, University of Chicago.

Abstract

The existence of advertising beyond the date at which a new good is introduced is anomalous in the context of Nelson-type models of advertising. This paper presents a model of advertising, based on private information about product quality, that i) yields such recurrent advertising in equilibrium; ii) generates equilibrium relationships among product price, advertising and quality similar to those suggested by earlier work, and iii) offers other new testable implications.

The formal structure is a two-period game of imperfect information, with the equilibrium concept being a refinement of perfect sequential equilibrium. An accessible introduction to the use and refinement of sequential equilibrium is provided; the power of such techniques for applied research is illustrated within the analysis of the advertising model.

I. INTRODUCTION

New goods are often advertised well beyond the time at which they are introduced. This observation is an elementary one, but one that is a major anomaly in the context of current theories of advertising. In existing models, after an initial advertising flurry, there is no value to continued advertising activity. This paper presents a simple two-period model of advertising behavior that gives such recurrent advertising a key role. It also generates a number of new hypotheses unique to the model while retaining propositions similar to the major results offered in previous work.

The basic structure of the model is straightforward. A firm has developed a new product and in the process of doing so has learned the good's "quality", which is private information. The firm may advertise as well as set an initial price at which it is willing to sell the good. Consumers draw what inferences they can from the firm's behavior and purchase if it pays for them to do so. Consumption experience is informative about quality, but only imperfectly so. Subsequently, the firm may advertise again and make a new price offer, and consumers may repurchase if they desire.

The principal manner in which the present model differs from earlier work [Nelson (1974), especially as formalized by Milgrom and Roberts (1986) and Kihlstrom/Riordan (1984)] is that a single consumption experience is not assumed to provide consumers with all the information to which the producer has access. Therefore, while consumption is informative, it is at best a noisy signal of quality. This assumption appears plausible. Consumption frequently involves a considerable amount of randomness either because of stochastic elements in the product itself or the circumstances in which consumption takes place. Moreover, it is transparent that something like this specification is necessary for recurrent advertising when the role of advertising is the transmission of information. If one consumption experience provides all available information, advertising after an initial sale is redundant (and, indeed, wasteful). In this sense, while it may at first appear to be a minor point, the assumption about the informativeness of consumption proves fundamental to the advertising model.

Section II sets out the model in detail and discusses the equilibrium concept. As is common in imperfect information environments, the model analyzed here has many sequential equilibria. The equilibrium concept used here is a slightly strengthened version of Grossman and Perry's (1986) perfect sequential equilibrium. Section II provides a discussion of the rationale behind such "refinements" of sequential equilibrium and explains how the one used here is employed.

Section III shows that there are two kinds of equilibria in this model. In one, if the firm has a good quality product to sell, it both advertises and offers a high price in both periods. All potential customers try the good initially, while only those who have a favorable consumption experience buy again. If the firm has a lower quality good to sell, it randomizes between the behavior of a producer of high-quality goods and charging a low price with no advertising. If the low-price/no-ad outcome occurs initially, quality is revealed in the process. In this case the same outcome necessarily occurs in the subsequent period. The important feature of this equilibrium--referred to as AD-AD--is that advertising can only occur at the outset if it might be recurrent.

In the other equilibrium, there is no advertising at the outset. Price is independent of quality and set such that all consumers are willing to buy initially. Subsequently, two outcomes are possible. In one, the product price falls relative to the introductory price, so that even consumers who had unfavorable consumption experience are willing to buy, and the good remains unadvertised. In the other, the subsequent behavior is exactly that occurring in the second period of the AD-AD equilibrium.

Section IV spells out the model's observable implications. The most important (from the point of view of a refutation) of these is that, should advertising occur after the good's introductory period, the product price must be higher than it was initially. This result holds in every equilibrium of the model and may be tested easily since it refers to the behavior of a single firm over time. In this case relatively little information is required.

A second set of empirical implications deals with outcomes that should be observed across firms. These focus on situations in which advertising is observed to occur when the good is introduced. This approach ensures that, if the model analyzed here is in fact generating the data, it is the AD-AD equilibrium that is doing so. It then follows that among other things, in the post-introductory period, Nelson-type relationships (e.g. advertising expenditures, product price and quality all being positively correlated) should be observed.

Finally, comparative statics experiments are considered. Some results are unusual. For example, again given advertising at the outset, an increase in production cost lowers the price of goods advertised outside the post-introductory period.

II. THE ADVERTISING GAME

The Problem

The situation under consideration is one in which a single firm is producing a good whose quality (q) can take on one of two values: h (high) or ℓ (low). These two outcomes might be thought of in terms of durability. For example, a high-quality battery functions longer than a low-quality battery. Alternatively, they might be interpreted as "quality of match", so that, within a given consumer population, high-quality products produce a larger fraction of "satisfied customers" than the low-quality product. In either case, if advertising is to serve any role at all, it must be that information on quality (i.e. the value of q) is distributed asymmetrically. In particular, it must be that the firm has better information on quality, at least initially, than do consumers. More importantly, for advertising to persist over time, it must be possible for this state of asymmetric information to persist as well.

This informational setup can be captured as follows. It is assumed that the good is such that its performance in the hands of consumers is stochastic. Specifically, the good either "works", in which case the consumer gets a utility normalized to 1 from a unit of consumption, or it "fails", providing the consumer a utility of $V \in (0,1)$. It is also assumed that in the process of developing, designing and testing the product, the firm has learned the probability that the product works (i.e. has learned quality). This probability is given by ψ_q $(q = h, \ell)$, with $\psi_\ell < \psi_h$. The consumer does not have access to this information initially. That is, while he knows the probability that the product he buys will work is one of ψ_h or ψ_ℓ, he does not know which it is (i.e. $q = h$ or $q = \ell$).

Because these assumptions differ from those generally adopted and are important in what follows, it is worth digressing somewhat to elaborate. First, it is clear that, under these assumptions, consumption experience alone is at best imperfectly informative. A single consumption experience provides the consumer only with the information that the product either worked or failed. Since it is assumed that the probability of the product working is correlated with quality, the outcome provides some information on quality. However, consumption experience by itself does not yield the same information on quality as the firm has obtained. For example, a single observation on a particular battery's durability yields only imperfect information concerning the entire distribution of death dates of batteries.

Second, consumption experience in a given period will be wholly uninformative about future consumption experience unless quality is correlated across periods. Additional complication is avoided if it is simply assumed that quality is fixed over time. In this way, consumption experience in the current period provides useful information about the quality of the good purchased in future periods.

Finally, unless the consumption experience of an individual is private information for that individual, mechanisms like warranties or money-back guarantees may eliminate advertising. As well, without private information, consumers could reliably use the consumption experience of others in making any repurchase decision. With private information, all such mechanisms become infeasible and advertising presents itself as a potential information transfer mechanism. Therefore, this model supposes that all consumption experience of an individual is private to that individual. The model, then, is applicable to products having the property that the characteristic(s) about which there is asymmetric information are also ones for which individual consumption experience is unverifiable.[1]

In addition to using consumption experience, consumers may also use observations on the outcomes of the firm decision-making process to draw inferences on quality.[2] The firm is assumed to make two choices: a unit price, p, for the good and a level of advertising, a. As usual, the choice of advertising level is a binary one involving either no ad (ϕ) or advertising at a level A, where the costs of these are given by zero and $\alpha > 0$ respectively. In each case, the choice should be thought of as a decision to advertise (or not) above the level necessary simply to inform customers of the product's existence, price, etc.

This latter point is an important one. The feature of the action called "advertising" that makes it a potentially useful signal is that the immediate effect is evidently costly to the firm on net (i.e. the immediate costs outweigh the immediate gains). Therefore what is important about advertising, from a signalling standpoint, is that part of the advertising expenditure in excess of the amount needed simply to inform consumers of the product's existence. Milgrom and Roberts make a similar point. Therefore, it is this part of the overall promotion activity that is defined as advertistment. This notion of advertising means that many things not commonly considered advertising in the everyday sense of the word, like a firm's support for the arts or amateur sports, might usefully be included in the definition of advertising. It also provides a simple explanation for the observed characteristics of much of what is commonly referred to as "ads".[3]

The problem that the firm faces is one of choosing a sequence of prices and advertising levels (one pair per period), given the quality of its product, in order to maximize the expected value of its profit stream. The firm makes these decisions realizing that consumers gather information both from the observed price/ad outcome and consumption experience. Consumers are faced with the problem of trying to infer from all their experience (both consumption and price/ad observations) the quality of the firm's product and, then, to make an expected utility-maximizing consumption decision. For simplicity, the consumer's decision is assumed to involve buying either zero or one unit of the product each period and all consumers are risk neutral (so that an individual buys if he obtains non-negative expected utility when prices are netted out). The complete structure of the consumer-firm game and a discussion of equilibrium is provided in the subsection following.

The Game Representation

The situation described above can be represented by an extensive form game with imperfect information. The main elements of this game and the necessary notation are provided below. In the interest of brevity, and to avoid the associated hailstorm of notation, some of the fine details of the game have been suppressed.

There are two kinds of active players: the firm (or producer), $\wp$, and a continuum of consumers, uniformly distributed on [0,N]. Any particular consumer will be labelled C, with $C \in [0,N]$.

Time is indexed by t with t=0,1 representing the initial and subsequent periods respectively. The quality (q) of $\wp$'s good is determined exogenously at t = 0 and, as mentioned above, is fixed for both periods. The outcome q = h is assumed to occur with probability $\delta \in (0,1)$, which also gives the consumer's initial priors on the outcome q = h. δ is the outcome of (unmodelled) costly and imperfect information gathering by C.[4] A firm having obtained the outcome q will be referred to as being of "type q";

$q = h, \ell$. Quality is assumed to be such that the unit cost of production, γ, is constant and independent of firm type.[5]

$\wp$'s strategy at $t = 0$ consists of a choice of a price offer for period 0, p_0, and an advertising choice, $a_0 = A$ or ϕ, the choice of each potentially depending on $\wp$'s type. To allow for randomization possibilities, a typical strategy for $\wp$ is represented by the function $F_0(p_0, a_0 | q)$, giving the probability that a price less than or equal to p_0 along with the ad choice a_0 is observed, given $\wp$ is of type q.

C's strategy at $t = 0$ consists of a decision either to buy a single unit of the good or not given the observation (p_0, a_0) and the prior probability on q, δ. For simplicity, C is restricted to pure strategy choice only. Should C purchase, consumption experience is as outlined previously. In particular, C's experience, c, may be either good (g) or bad (b) with $Pr(c=g|q=h) = \psi_h > Pr(c=g|q=\ell) = \psi_\ell$. An independent realization of either c=g or b is obtained by each C at t=0 and this realization is private information (as discussed previously).

In period t=1 the situation is similar to that in t=0 with two exceptions. One is that some (or indeed all) C may not have purchased at t=0 while others may have. It is easy to show, however, that there could not be an equilibrium in which $\wp$'s strategy results in some C not purchasing at t=0. Therefore, to simplify the exposition, the remainder of this section focuses only on the situation in which all C purchase at t=0. The other difference is that, at t=1, $\wp$'s strategy may specify an action (p_1, a_1) that results in only those C having had a good consumption experience in t=0 purchasing, only those having had a bad experience purchasing, or all C purchasing. Of course, while $\wp$ cannot observe who these various individuals are, there is enough information to calculate the expected number in each case (i.e. $\psi_q N$, $(1-\psi_q)N$, or N).

Otherwise, the specification of the t=1 problem is entirely analogous to that at t=0. A typical strategy for φ is given by the function $F_1(p_1,a_1|q,p_0,a_0)$ and is interpreted as before with the addition that F_1 may depend on p_0 and a_0 as well as q. For brevity's sake, this dependence on p_0 and a_0 will be implicit subsequently. *C*'s strategy, again, consists of a purchase or no purchase decision and may be conditioned on p_0,a_0,p_1,a_1 and consumption experience c.

The description of the game is completed by a specification of the payoffs. *C*'s payoffs need not be presented in any detail. Conveniently, the recursion relationship studied below yields the necessary information as a by-product.

As for φ's payoff, consider, first, the expected profit earned from sales at t=1 given the strategy pair F_0,F_1. Since *C*'s strategy specifies a purchase/no purchase decision for every (p_1,a_1), it is possible to specify an acceptance set, B_1, defined as the set of (p_1,a_1) such that *C* purchases at t=1. This set is simply all (p_1,a_1) yielding *C* non-negative expected utility. It therefore depends on the prior δ, the strategies F_0,F_1 and *C*'s consumption experience at t=0. If all *C* are treated symmetrically, in the sense that each draws identical inferences from the observation of the (p_t,a_t) pairs regardless of consumption experience, then the acceptance set for those *C* with bad experience at t=0, $B_1(b)$, is a subset of the one for those with good experience, $B_1(g)$.[6]

It follows that φ's expected profit from sales at t=1 for some strategy pair F_0,F_1 is

$$\pi(p_0,a_0|q) = \int_{B_1(g)-B_1(b)} [N\psi_q(p_1-\gamma) - \alpha(a_1)]dF(p_1,a_1|q)$$

$$+ \int_{B_1(b)} [N(p_1-\gamma) - \alpha(a_1)]dF(p_1,a_1|q)$$

where $\alpha(a_1) = 0$ if $a_1 = \phi$ and α if $a_1 = A$. The first term in this expression is the expected return to φ should it choose a (p_1,a_1) that results in a purchase only by those *C* with a good consumption experience at t=0. (i.e. (p_1,a_1) in the set $B_1(g) - B_1(b)$). Any one of these (p_1,a_1) is chosen with "probability" $dF(p_1,a_1|q)$. Since only $\psi_q N$ of the *C* are expected to have good experience, variable returns are proportional to that number. The second term is the expected return to φ should it choose a (p_1,a_1) that results in those *C* with bad experience in t=0 purchasing. Since those with good experience purchase if those with bad experience do, variable returns here are proportional to N.

φ's profits from sales at t=0 are defined similarly, with *C*'s strategy determining an acceptance set B_0 giving the (p_0,a_0) that result in a purchase by *C* at t=0. The only difference is that, now, *C* may purchase at t=0 even though expected utility in t=0 is negative. That is, *C* will purchase as long as the expected utility stream from such a decision were higher than that from not purchasing in t=0 and then proceeding optimally.

φ's payoff -- total expected profit -- is then defined as (ignoring discounting)

$$\pi(q) = \int_{B_0} [N(p_0-\gamma) - \alpha(a_0) + \pi(p_0,a_0|q)]dF_0(p_0,a_0|q). \qquad (1)$$

That is, given (p_0, a_0), $\wp$ earns $(p_0-\gamma)$ for each of the N units sold, and spends $\alpha(a_0)$ on advertising. This outcome is followed by the period 1 profit $\pi(p_0, a_0|q)$ defined previously. Under the strategy $F_0(\ )$, (p_0, a_0) occurs with probability $dF(p_0, a_0|q)$.

Additional Restrictions

It is possible to analyze the game set out above without further restrictions. However, a variety of uninteresting outcomes can be eliminated through the imposition of three parameter restrictions. Because these restrictions all involve the one period expected utility for *C*, it is useful to define the variable $\mu_q \equiv \psi_q + (1-\psi_q)V$. μ_q is *C*'s expected utility from consumption, given that quality is q, and provides information about *C*'s demand behavior (acceptance set). For instance, no *C* could credibly claim to be unwilling to purchase a unit of the good at a price μ_ℓ or less. Similarly, at t=1, no *C* could credibly claim to be willing to purchase a unit of the good at a price beyond μ_h.

Given these characteristics of *C*'s demand behavior, one obvious restriction is:

$$\mu_\ell - \gamma \geq 0 \tag{2}$$

In the absence of (2), no action that reveals $q = \ell$ can ever be part of an equilibrium. For were $\wp$ of type ℓ, ceasing production would be preferable to the claimed equilibrium action. Therefore, (2) is required to make separation feasible, at least.

The next two restrictions are less obvious but needed to guarantee that advertising is at least possible in each period. The restrictions are:

$$N(\mu_h - \gamma)\psi_\ell - \alpha > N(\mu_\ell - \gamma) \tag{3}$$

and

$$\alpha < N(\mu_h - \mu_\ell) \tag{4}$$

Equation (3) guarantees the existence of a p_1 with the property that a φ of type ℓ would prefer to advertise and set price p_1, selling only to those C having good experience in t=0, to revealing itself as type ℓ and selling to all C. If (3) does not hold, then, there can never be sufficient return to a φ of type ℓ to induce it to advertise in t=1. φ would simply prefer to reveal itself as type ℓ and sell to all C. Thus, (3) allows for the possibility that a φ of type ℓ may be mimicking a φ of type h by advertising at t=1.

Equation (4) has a similar interpretation. If (4) is violated, then p_0 cannot possibly be set high enough to make the return to advertising at t=0 dominate simply not advertising and selling to all C at $p_0 = \mu_\ell$. Of course, φ might expect to recoup some of these early losses by advertising at t=1. However, as will be seen subsequently, this possibility will prove infeasible in equilibrium.

Definition of Equilibrium

In the next section, the equilibria of the game set out above are constructed. It is shown that the game has at most two kinds of equilibria satisfying a strengthened version of Grossman and Perry's perfect sequential equilibrium (PSE) notion. This subsection discusses the nature of the refinements employed and tries to provide some justification for their use. The reader who finds such discussions objectionable may wish to skip this subsection and merely interpret the equilibria discussed below as two of a large set of Nash (or, if it is preferred, sequential) equilibria.

A Nash equilibrium (NE) of the ad game is a pair of functions, $F_0(p_0,a_0|q)$ and $F_1(p_1,a_1|p_0,a_0,q)$, and a purchase rule by *C* such that: i) given $F_0(\cdot)$ and $F_1(\cdot)$, no *C* can earn a higher expected utility by altering the purchase rule (i.e. *C*'s strategy must be a best reply to F_0,F_1), and ii) given *C*'s purchase rule, $\wp$ can earn no higher expected payoff by altering F_0 or F_1. As is now well known, games such as this one may have a large number of Nash equilibria, many supported by incredible threats by one (or both) of the players. For example, in this game, even though $\mu_\ell > \gamma$, the outcome $a_t \equiv \phi$, $p_t \equiv \gamma$ can result in a Nash equilibrium. This outcome is supported by *C* refusing to purchase at any $p_t > \gamma$ (independent of a_t).

The approach adopted to deal with such outcomes is to use the sequential equilibrium (SE) concept of Kreps and Wilson (1982). Basically, it is required that each player's actions be ones that are in his interest to undertake if he is ever called upon to do so. In the proposed equilibrium above, for instance, *C* would not be permitted to adopt the action "no purchase" if he observed a p_t such that $\gamma < p_t < \mu_\ell$ because purchase is sure to yield positive expected utility. Therefore, the above Nash equilibrium would not be a sequential equilibrium.

In many circumstances, a player is required to specify an action as part of a sequential equilibrium when he is uncertain about the payoff that the action will produce. *C*, for instance, may have to make a purchase decision while still uncertain as to the outcome, q. In order to be sure that the action specified is one that is in a player's interest to carry out ("sequential rationality") sequential equilibrium requires that a player's beliefs about his position in the game be specified, so that expected utility may be computed. In the current context, doing so requires that functions $\rho_0(p_0,a_0),\rho_1(p_0,a_0,p_1,a_1,c)$ be specified, defining C's beliefs about the outcome q as a function of all observables. The function $\rho_0(\cdot)$, for instance, specifies the probability C attaches to the outcome $q = h$ prior to the purchase at t=0 given observation of any (p_0,a_0). $\rho_1(\cdot)$ is interpreted similarly for t=1, with c being the consumption experience in t=0 and taking either of the values b or g. A sequential equilibrium requires that, for (p_0,a_0) and (p_1,a_1) that are the outcome of equilibrium play, $\rho_0(\cdot)$ and $\rho_1(\cdot)$ be constructed from the equilibrium strategies F_0, F_1 and δ using Bayesian updating. Of course, since players must specify actions that they would be willing to take but may in fact not have to take in the equilibrium, beliefs must also be specified for (p_0,a_0) and (p_1,a_1) that may not be observed outcomes of the equilibrium play. In these situations, players' beliefs are effectively unrestricted.

While Nash equilibria supported by incredible threats are ruled out by the sequential equilibrium concept, many equilibria still remain that meet the sequential equilibrium requirements. While none of these can be supported by incredible threats, many are supported by what has come to be described as "incredible beliefs". To illustrate, the outcome $a_t \equiv \phi$, $p_t \equiv \mu_\ell$ can be the result of a sequential equilibrium in which C refuses to purchase at any price above μ_ℓ (independent of a_t). This situation is supported by C believing that, if $p_t = \mu_\ell$, $a_t = \phi$, then $q = h$ with probability δ, while for any other observed (p_t,a_t) pair, $q = \ell$ with certainty. With these beliefs, it is in C's interest to buy only if $p_t \leq \mu_\ell$ and, with this strategy for C, it is in $\wp$'s interest, regardless of type, to set $p_t = \mu_\ell$, $a_t = \phi$. Further, when $p_t = \mu_\ell$, $a_t = \phi$ is observed, beliefs are consistent with $\wp$'s equilibrium strategy and Bayesian updating of the prior δ.

In this case, the threat by C not to buy if $p_t > \mu_\ell$ is credible because, if $p_t > \mu_\ell$ is believed to signal $q = \ell$ for sure, it is not in C's interest to purchase. However, those beliefs may be incredible in a way analogous to the manner in which certain threats are incredible. In the case of incredible threats, if called upon to carry out the threatened action, the player could do better by acting in some other fashion. Similarly, if some action is observed, the player must indeed interpret that action in the way specified by his beliefs. In the example, C must interpret any $p_t > \mu_\ell$ as implying $q=\ell$. If, having observed some $p_t > \mu_\ell$ that is not an equilibrium outcome, there is no reason to think $q = \ell$ necessarily holds, then believing $q=\ell$ is incredible. It could be that $p_t > \mu_\ell$ may equally have been set by either type firm.

If the alternative beliefs (i.e. $p_t > \mu_\ell$ equally likely set by either type) are adopted, however, the outcome $p_t = \mu_\ell$, $a_t = \phi$ cannot be a sequential equilibrium. At t=0, either type φ could set $p_0 = \delta\mu_h + (1-\delta)\mu_\ell$, $a_0 = \phi$, and increase profits. The increased profits result from the fact that, if C holds the beliefs that $p_0 > \mu_\ell$ is equally likely set by a type ℓ as a type h firm, it is in C's interest to purchase when $p_0 = \delta\mu_h + (1-\delta)\mu_\ell$. Thus $p_t = \mu_\ell$, $a_t = \phi$ can arise as the outcome of a sequential equilibrium only because of the incredible beliefs employed by C.

As sequential equilibrium reduces the set of Nash equilibria by ruling out incredible threats, so many "equilibrium refinements" have been proposed to reduce the set of sequential equilibria (and so NE) by ruling out incredible beliefs [see Grossman and Perry (1986), Cho and Kreps (1987), Banks and Sobel (1987), Farrell (1985)] in a structured manner. The refinement adopted here is a variant of Grossman and Perry's perfect sequential equilibrium (PSE). To understand how the equilibrium concept employed here works, it is necessary to understand first what a PSE is in the context of the ad game.

The aspect that distinguishes a PSE from an SE is embodied in the nature of a credible updating rule. To see how this concept operates, consider a strategy for $\wp$ at t=1, $F_1(p_1,a_1|q)$, part of an SE, and proposed as part of a PSE. (Since PSE is a refinment of SE, any PSE must be an SE.) Now, let $\wp$ consider an alternative strategy $G(p_1,a_1|q)$. $G(\cdot)$ is called a "signalling strategy". It may involve $G \equiv F_1$ for some types, q, and may be such that the support of G intersects that of F_1 for every q. However, $G(\cdot)$ must allow some (p_1,a_1) not possible under F_1. Also, assume that in assessing $G(\cdot)$, $\wp$ imagines that for any (p_1,a_1) in the intersection of $G(\cdot)$ and $F_1(\cdot)$'s supports, C's beliefs are as specified by the SE beliefs constructed from F_1. For (p_1,a_1) in the support of G but not F_1, C's beliefs are constructed from G using Bayes' rule. These beliefs may be different than those in the SE of which $F_1(\cdot)$ is a part. (Beliefs attached to other (p_1,a_1) may be ignored for this discussion.) Suppose, also, that C's actions, given these beliefs, are maximal for C (i.e. in C's interest to carry out).

Given this setup, assume $G(\cdot)$ to have the properties that i) any randomization for some type q involves only (p_1,a_1) that are equally good for that type; ii) when $G(\cdot|q) \equiv F_1(\cdot|q)$, $\wp$ is at least as well off following F_1 as G for that q; iii) for some q, following $G(\cdot)$ is strictly preferred to following F_1. Then, a credible updating rule for the strategy pair F_0, F_1, must specify that beliefs adopted by C for (p_1,a_1) in the support of G but not F_1 be consistent with $G(\cdot)$ and Bayesian updating. If there is more than one $G(\cdot)$ that satisfies this requirement, then credible updating allows a choice from among the $G(\cdot)$. If there are no $G(\cdot)$ that satisfies the requirement, then credible updating implies no restriction beyond those given by the SE.

The impact of this credibility restriction comes from the fact that the beliefs implied by credible updating for those (p_1, a_1) that will not occur, given F_1, may be incompatible with F_1 being part of an equilibrium. To illustrate, suppose that the ad game were a single-period game, so that $\wp$'s strategy would be a function $F(p,a|q)$ and C's information prior to observing (p,a) simply that $Pr(q = h) = \delta$. Then, as previously, $(p,a) = (\mu_\ell, \phi)$ occurring with certainty is a sequential equilibrium of this game, supported by the beliefs for C that $\rho(p,a) = 0$ unless $(p,a) = (\mu_\ell, \phi)$. However, it is not a PSE. To see why, note that any $(\hat{p}, \phi)$ with $\hat{p} \in (\mu_\ell, \delta\mu_h + (1-\delta)\mu_\ell]$ earns $\wp$ higher profits for any q, should C purchase, than does (μ_ℓ, ϕ). Further, since $\delta\mu_h + (1-\delta)\mu_\ell$ is C's expected utility given only the prior information, C will purchase given $(\hat{p}, \phi)$ if $\rho(\hat{p}, \phi) = \delta$. Therefore, a signalling strategy always open to $\wp$ is to choose a $G(\cdot)$ such that both types pick $(\hat{p}, \phi)$ with probability one. Bayesian updating will imply that $\rho(\hat{p}, \phi) = \delta$ and both types of $\wp$ will be better off. Thus $(p,a) = (\mu_\ell, \phi)$, supported by the beliefs given above, cannot be part of a PSE. Indeed, it turns out that there is no signalling strategy that can result in $\rho(\hat{p}, \phi) = 0$ for all $\hat{p} \in (\mu_\ell, \delta\mu_\ell + (1-\delta)\mu_\ell]$. Thus, credible updating implies the outcome (μ_ℓ, ϕ) cannot be a part of any PSE.

The equilibrium concept used below is a further refinement of PSE. The additional restrictions arise from the fact that the ad game lasts more than one period and that the equilibrium outcomes at t=1, conditional on t=0 outcomes, are often not unique.

The first restriction relates to the way that *C*'s beliefs at t=1 can be conditioned on t=0 outcomes. At the beginning of t=1, *C* has observed (p_0, a_0) and obtained consumption experience c = g or b. At this point, *C* can construct a belief $\rho_c(p_0, a_0, c)$ from $\rho_0(p_0, a_0)$ and the outcome c using Bayesian updating. The restriction that is imposed on second-period beliefs $\rho_1(\cdot)$ is that $\rho_1(p_1, a_1, p_0, a_0, c)$ can be written as $\rho_1[p_1, a_1, \rho_c(p_0, a_0, c)]$. This restriction simply requires that any first-period outcomes that result in the same information for *C* at t=1, must yield the same beliefs at t=1 when some outcome (p_1, a_1) is observed. This restriction prevents outcomes that have identical informational consequences from being followed by drastically different equilibrium outcomes at t=1.

The second restriction is in the nature of a continuity restriction. It arises because of the fact that, even with the above restriction, there are generally multiple PSE strategies for φ at t=1 given any $\rho_c(p_0, a_0, c)$. The beliefs that support each are, as one might expect, quite different, as are the $F_1(\cdot)$ themselves. This latter fact provides the motivation for the particular form that the restrictions take. Indeed, the impact of the restrictions is to require that if period 0 actions (p_0, a_0) are followed by a period 1 outcome having certain characteristics, then a signalling deviation, yielding information almost identical to that provided by (p_0, a_0), both can and will be followed by period 1 outcome having nearly the same characteristics as that following (p_0, a_0).

To understand the restriction more fully, consider the issue of whether some strategy pair $\tilde{F}_0, \tilde{F}_1$ is part of an equilibrium. Signalling strategies,

G_0 followed by some F_1 that could be supported as part of a PSE for t=1, are examined, as above, to check whether the updating rule in use is credible. PSE reasoning supports $\tilde{F}_0$ and $\tilde{F}_1$ as long as there is <u>one</u> credible updating rule from which beliefs are constructed. Such a rule might <u>require</u> (i.e. no other rule will support $\tilde{F}_0, \tilde{F}_1$ as a PSE) that a signalling strategy G_0, yielding a value $\rho_c(p_0, a_0, c)$ arbitrarily close to the $\rho_c(\cdot)$ implied by $\tilde{F}_0$, be followed by a PSE involving F_1 with implied beliefs very different from those in $\tilde{F}_1$. This outcome is possible because of the multiple PSE in t=1. The continuity restriction would make such updating rules inadmissable. In formal terms, the updating rule is a function λ, associating with any value of ρ_c, a system of beliefs $\rho_1(p_1, a_1, \rho_c)$. The restriction is that it be possible to support the equilibrium with an updating rule λ that is continuous for (at least) ρ_c in a neighborhood of the value of ρ implied by $\tilde{F}_0$. An equilibrium supported by an updating rule obeying this and the preceding restriction will be referred to as an MPSEC (Markov Perfect Sequential Equilibrium with Continuous Updating).

The intuition behind these restrictions is quite simple. Consider an action by $\mathfrak{p}$ at t=0 that results in only a small change in C's assessment of quality at t=1. The restrictions require that such an action not result in a large change in C's purchase decision at t=1. In particular, if there is some p_1 at which C was previously willing to purchase in equilibrium, there must still be p_1 near by at which C is willing to purchase in the new equilibrium. Again, if the reader views these requirements as too restrictive, it should be recalled that the equilibria developed below are part of a larger class of equilibria (either NE, SE, or PSE). The ones considered can be viewed as particularly interesting elements of this class.

III. EQUILIBRIUM IN THE AD GAME

In this section the MPSEC of the ad game are constructed. The procedure is as follows: A period 0 outcome (p_0,a_0) generates some initial assessment by *C* of $\Pr(q = h|p_0,a_0)$, $\rho(p_0,a_0)$. This value, combined with his consumption experience, c, produces a period 1 assessment for *C*, $\Pr[q = h|p_0,a_0,c] = \rho_c(p_0,a_0,c)$. Since t=0 outcomes affect *C*'s decisions only through $\rho_c(\cdot)$ the first step must be to consider the outcomes that can be supported as the t=1 portion of a MPSEC for different values of $\rho_c(p_0,a_0,c)$. Having done so, the period 0 problem can be considered [i.e. what (p_0,a_0) outcomes can occur in a MPSEC].

1. Period 1

A useful way to begin is to determine what *C* would be willing to pay in period 1 purely on the basis of information obtained in period 0 as summarized by $\rho_c(\cdot)$. Recall that *C* obtains expected utility of μ_q, given the good is of quality q. Because period 1 is the final consumption period, given no new information, *C* is willing to pay at most

$$\rho_c(p_0,a_0,c)\mu_h + [1 - \rho_c(p_0,a_0,c)]\mu_\ell .$$

When $\rho(p_0,a_0) = 1$ or 0, $\rho(p_0,a_0) = \rho_c(p_0,a_0,c)$ for all c. In this case any C is willing to pay μ_h (μ_ℓ) when $\rho(p_0,a_0)$ takes on the value 1(0). Otherwise, consumption experience provides additional information and $\rho_c \in (0,1)$. For those C having experience $c = g$ ($c = \ell$), define p_H (p_L) as the highest price they would be willing to pay without further information. Sequential rationality implies that, should C receive no additional information in t=1, he cannot claim to be unwilling to purchase at some price less than his maximum willingess to pay (i.e. p_H for $c = g$, p_L for $c = \ell$).

Now consider $\wp$'s strategy should period 0 somehow provide perfect information: $\rho(p_0,a_0) = 1$ or 0. When $\rho(\cdot) = 1$, any C is willing to pay μ_h. It follows that

$$F_1(p_1,a_1|q) = \begin{cases} 1 & p_1 \geq \mu_h,\ a_1 = \phi \\ 0 & \text{otherwise;} \end{cases}$$

i.e. $(p_1,a_1) = (\mu_h,\phi)$ holds with certainty, irrespective of q. The argument is that any $p_1 > \mu_h$ yields no sale, while any $p_1 < \mu_h$ could be improved upon by a deviation by either type $\wp$ to $p = \mu_h$.[7] Any nontrivial advertising could also be improved upon in a similar fashion. Analagously, when $\rho(\cdot) = 0$, $(p_1,a_1) = (\mu_\ell,\phi)$ must occur, with certainty.

When period 0 actions are not fully informative about quality, the situation becomes more complicated. Nevertheless, it is easy to show that there are three kinds of potential equilibrium outcomes:

a) "High Pooling" (HP) -- $F_1(p_1,a_1|q) = \begin{cases} 1 & p_1 \geq p_H,\ a_1 = \phi \\ 0 & \text{otherwise} \end{cases}$

That is, independently of q, 𝔭 charges the maximum willingness to pay for those *C* who had favorable consumption experience in t=0 and 𝔭 does not advertise.

b) "Low Pooling" (LP) -- $F_1(p_1,a_1|q) = \begin{cases} 1 & p_1 \geq p_L,\ a_1 = \phi \\ 0 & \text{otherwise.} \end{cases}$

When LP occurs, 𝔭 invariably offers a price equal to the maximum willingness to pay of those *C* having bad experience in t=0. 𝔭, therefore, sells to all *C*. Again, no advertising occurs.

c) "Advertising" (AD) --

$$F_1(p_1,a_1|q) = \begin{cases} 1 & p_1 \geq \hat{p},\ a_1 = A,\ q = h \\ 0 & p_1 \leq \hat{p},\ \text{or } a_1 = \phi;\ q = h \\ 1 & p_1 \geq \hat{p},\ a_1 = A,\ q = \ell \\ f & \mu_\ell \leq p_1 < \hat{p},\ a = \phi,\ q = \ell \\ 0 & \text{otherwise,} \end{cases}$$

for some $f \in (0,1)$. That is, for type h, 𝔭 charges a price $\hat{p}$ and always advertises. For type ℓ, 𝔭 randomizes between type h's actions and revealing his type by charging $p_1 = \mu_\ell$ and failing to advertise, this latter pair occurring with probability f. Only those *C* who had good consumption experience, c = g, purchase given $(p_1,a_1) = (\hat{p},A)$. All *C* would purchase having observed $(p_1, a_1) = (\mu_\ell, \phi)$. Because $(\hat{p}, A)$ is more likely to be chosen if $\wp$ is of type h than of type ℓ, the observation of this pair provides additional information to C. How likely C believes it is that q=h, given an observation of $(\hat{p}, A)$, depends on the probability with which $\wp$ of type ℓ chooses the pair. Therefore, a restriction on f is that it be sufficiently large (1-f small) that those C with good experience at t=0 are willing to purchase at $\hat{p}$. Also, $\hat{p}$ must be chosen so that $\wp$ is willing to randomize when $q = \ell$, which requires that

$$N(\hat{p} - \gamma)\psi_\ell - \alpha = N(\mu_\ell - \gamma),$$

or $$\hat{p} = \gamma + \frac{\mu_\ell - \gamma + \frac{\alpha}{N}}{\psi_\ell}.$$

Note that (4) guarantees that $\hat{p} < \mu_h$, so that some $f < 1$ will succeed in generating a purchase when $(p_1, a_1) = (\hat{p}, A)$.

The demonstration that if there are any MPSEC, the t=1 portion must be of the form HP, LP or AD, is not difficult. First, consider possible pure strategy equilibria; that is, equilibria in which $F_1(p_1, a_1|q)$ yields some (q-dependent) pair (p^q, a^q) in t=1 with certainty, for each q. Suppose that the equilibrium is "separating"; that is $(p^h, a^h) \neq (p^\ell, a^\ell,)$, so that q may be inferred from (p_1, a_1). If beliefs are consistent with such an equilibrium outcome, it must be that either all C purchase or no C purchase, regardless of consumption experience, which means that, for this outcome to occur, $\wp$'s payoff must be independent of q. Otherwise, it would pay $\wp$ to deviate to whichever of (p^q, a^q) offered the higher payoff. Further, since no C would be willing to pay more than μ_ℓ for the good, having observed (p^ℓ, a^ℓ), it must be that the common payoff is no larger than $N(\mu_\ell - c)$; that is, the payoff a type ℓ earns by setting $(p_1,a_1) = (\mu_\ell,\phi)$.

Consider, however, a signalling strategy, G, in which $(p_1,a_1) = (p_L,\phi)$ occurs with certainty for each q. Such a strategy yields φ a payoff strictly greater than that attained in the proposed equilibrium, irrespective of q. This outcome occurs because no customers are lost (due to C being required to have beliefs, having observed (p_L,ϕ), consistent with $G(\cdot)$ and $p_L > \mu_\ell$ if $\rho_c \in (0,1)$. Thus, beliefs that would result in $(p^h,a^h) \neq (p^\ell,a^\ell)$ as an equilibrium outcome cannot be consistent with credible updating, which means that only outcomes with $(p^h,a^h) = (p^\ell,a^\ell)$ can be part of a MPSEC that specifies a pure strategy in t=1.

Given the absence of separation, it follows that a pure strategy F_1 can only yield HP or LP. Any outcome involving advertising is unsupportable due to the existence of the signalling strategy in which the same price is charged but no advertising is undertaken. Any price below p_H (p_L) cannot be supported because of the availability of the signalling strategy in which p_H (p_L) is charged. Whether (p_H,ϕ) or (p_L,ϕ) is an equilibrium outcome depends on a variety of factors, and is explored subsequently.

Next, consider strategies specifying nontrivial randomization in period 1. Since all the points over which randomization is to occur must yield the same payoff, given q, it is immediate that randomization for a given type can occur over at most four points. Of these four points, two can be such that all C purchase should the point be observed and two such that only those C for whom c = g purchase. In addition, two of the four can have advertising; however, these two cannot both be such that all C (only those with c = g) purchase.

It is also simple to show that, in any equilibrium, type h cannot randomize over any points that are not possibly chosen by type ℓ as well. To see this, suppose, to the contrary, that some equilibrium involved $\wp$ adopting such a strategy. The discussion above regarding the separating equilibrium implies that any such strategy must of necessity have the property that any point chosen by ℓ with positive probability is also chosen by h with positive probability. This observation immediately implies that $\wp$'s t=1 payoff for type h under such a strategy must equal or exceed the payoff for type ℓ. Should the payoff for type h exceed that for type ℓ, however, type ℓ could deviate to the point chosen only for type h and achieve the same payoff as $\wp$ does for type h. This possibility exists because beliefs consistent with the proposed strategy being an equilibrium would result in all *C* purchasing, regardless of consumption experience, should this point be observed. Therefore, the proposed strategy could only be an equilibrium should $\wp$'s t=1 payoffs be independent of q. For this outcome to occur, it must be that the points that $\wp$ sets with positive probability for both types result in all *C* purchasing. Further, one of these points must be more likely observed when q = ℓ than when q = h. This feature, however, means that a signalling strategy specifying (p_L,ϕ) with certainty yields a higher payoff for both types, making the beliefs supporting the proposed equilibrium strategy not form a credible updating rule. Therefore, it cannot be part of a MPSEC.

Next, it must also be that only one type randomizes in any equilibrium. Should an equilibrium involve $\wp$ randomizing for both types then either all points result in the same purchase behavior for *C* (in which case the randomization is over the two points at which all *C* purchase or the two at which those *C* for whom c = g purchase) or some points result in all *C* purchasing and others only those *C* having c = g purchasing. In the former case, the payoff to type q is bounded above by $(p_L - \gamma)N-\alpha$ $((p_H-\gamma)\psi_q N-\alpha)$. A signalling strategy specifying (p_L,ϕ) $((p_H,\phi))$ with certainty would yield a larger payoff to both types and, therefore, the proposed equilibrium could be supported with beliefs generated by credible updating. In the latter case, the (p_1,a_1) resulting in all *C* purchasing yield the same t=1 payoffs to both types. Those (p_1,a_1) resulting in purchase only by *C* for whom c=q yield a higher payoff for type h than for type ℓ (since $\psi_h > \psi_\ell$). Therefore, $\wp$ must be unwilling to randomize for one type in this case.

Finally, in any equilibrium, $\wp$ can only randomize when $q = \ell$. For were the randomization to involve h, instead, some (p_1,a_1) pair chosen by both types with positive probability would signal type ℓ (i.e. it would occur with higher probability when $q = \ell$ than $q = h$). In this case, a signalling strategy specifying one of (p_H,ϕ) or (p_L,ϕ) with certainty would yield a higher payoff to both types. This, again, implies that the beliefs supporting the proposed equilibrium are not from a credible updating rule so that randomization by type h is not part of a MPSEC.

Given this result, it only remains to determine the point chosen by h and the point(s), other than this one, over which ℓ randomizes. As for the latter, since any point chosen only by ℓ reveals quality as $q = \ell$, ℓ's best option is to choose only the point (μ_ℓ,ϕ). Any other point must yield a t=1 payoff lower than this point and so could be improved upon by a deviation. The (p_1,a_1) chosen with positive probability by $\wp$ for both types will be $(\hat{p},A)$ with $\hat{p}$ solving

$$N(\hat{p} - \gamma)\psi_{\ell} - \alpha = \mu_{\ell} - c,$$

where this equation guarantees that $\wp$ is willing to randomize between $(\hat{p},A)$ and (μ_{ℓ},ϕ) for $q = \ell$. To see this, note that the only other alternative (given a mixed strategy equilibrium exists at all) is $(p_1,a_1) = (\tilde{p},\phi)$ with $\tilde{p}$ the solution to

$$N(\tilde{p} - \gamma)\psi_{\ell} = \mu_{\ell} - c.$$

However, a signalling strategy involving exactly the behavior given for the AD outcome results in type ℓ being no worse off and type h strictly better off, implying that randomization over $(\tilde{p},\phi)$ and (μ_{ℓ},ϕ) cannot be supported by a credible updating rule.[8] Thus, if randomization occurs in period 1, advertising must be involved should beliefs be constructed in a credible fashion. This outcome follows from the fact that advertising is profitable for $\wp$. Updating rules supporting any other randomizations do so by constraining *C* to ignore this fact even should he be presented with evidence to support it. As such, the rules commit *C* to ignoring the value of advertising as a signal of quality.

To summarize, then, there is a MPSEC in which $\wp$ randomizes in period 1. The strategy must specify that $\wp$ randomize over (μ_{ℓ},ϕ) and $(\hat{p},A)$, the latter being chosen with certainty when $q = h$. The randomization is such that only those *C* for which $c = g$ purchase when $(\hat{p},A)$ is observed.

To demonstrate that the HP, LP and AD outcomes may in fact be part of a MPSEC is much more tedious, although no more complicated, than to demonstrate that there are no other possibilities. Now, it must be checked that no signalling strategy exists that will undermine the credibility of the updating rule supporting these outcomes. There are, of course, many candiates for such signalling strategies. Nevertheless, it is possible to verify that under some conditions each is part of a MPSEC.

The conditions under which the different possible equilibria arise are summarized in Table 1. There are five restrictions, all listed at the bottom of the table. Because these restrictions are not independent, they define a relatively small number of cases--10 to be exact. These 10 cases are exhaustive. In the first 6 cases only, a pure strategy outcome (HP and/or LP) may occur. In all cases, the AD outcome is possible if an additional restriction is satisfied. Should this restriction fail, then nonexistence occurs in cases 7-10.

All the restrictions serve to prevent certain kinds of signalling strategies from destroying credibility of the updating rule. For example, consider what is required to support HP in case 1. Several kinds of signalling strategies are possible. Pure strategy ones are either separating--in which case the signalling strategy specifies a separate (p_1,a_1) for each type--or "pooling"--wherein both types choose the same (p_1,a_1). A separating signalling strategy must result in purchase by all C regardless of consumption experience (or purchase by no C). As such, a strategy that is preferred for type h to the equilibrium strategy must also be preferred for type ℓ (obviously type ℓ will not choose a signalling strategy that reveals $q=\ell$.). However, this contradicts the fact that the signalling strategy

Table 1

Period 1 Outcomes

Sign of Restrictions[†]

Case Number	R1	R2	R3	R4	R5	Outcome*
1	+	+	+	+	+	HP,AD
2	-	-	~	~	~	LP,AD
3	-	+	-	~	~	LP,AD
4	-	+	+	-	-	LP,AD
5	+	+	+	-	+	HP,AD
6	-	+	+	-	+	LP,HP,AD
7	+	+	+	+	-	AD
8	+	-	~	~	~	AD
9	+	+	-	~	~	AD
10	+	+	+	-	-	AD

Restrictions

R1: $(\mu_h - \gamma)\psi_\ell - (p_L - \gamma)$

R2: $(p_H - \gamma)\psi_h - (p_L - \gamma)$

R3: $(p_H - \gamma)\psi_\ell - (\mu_\ell - \gamma)$

R4: $(p_H - \gamma)\psi_\ell - (p_L - \gamma)$

R5: $N(p_H - \gamma)\psi_\ell - [N(\mu_h - \gamma)\psi_\ell) - \alpha$

[†]Restrictions listed below Table. "+" denotes ≥ 0, "-" represents ≤ 0, and "~" is $\gtreqless 0$.

[*]All AD outcomes require the imposition of

$$N(\hat{p} - \gamma)\psi_h - \alpha \geq \max\{N(p_H - \gamma)\psi_h, N(p_L - \gamma)\}$$

is separating. A pooling strategy involving (p_1, a_1) either generates no sale $(p_1 > p_H)$, or less revenue for both types (immediate if $p_L < p_1 < p_H$, and implied by restrictions R2 and R4, if $p_1 \leq p_L$, irrespective of a_1).

Randomized signalling strategies are also possible. Again, $\wp$ cannot randomize for both types. Strategies in which $\wp$ randomizes for type h alone (and involve type ℓ choosing some point other than (p_H, ϕ)) contain a (p_1, a_1) outcome signalling $q = \ell$. This yields $\wp$ a payoff strictly less than that earned by following the equilibrium path, for $q = \ell$. When $\wp$ randomizes for $q = \ell$ only, there are two possibilities. If (p_H, ϕ) does not occur with positive probability in the signalling strategy, type ℓ's payoff is at most $N(\mu_\ell - \gamma)$ because some (p_1, a_1) reveals $q = \ell$. Since $N(\mu_\ell - \gamma) < N(p_L - \gamma) \leq N(P_H - \gamma)\psi_\ell$ by R4 ≥ 0 in case 1, type ℓ will not undertake the signalling strategy. If (p_H, ϕ) is part of $\wp$'s randomization for type ℓ, the pair $(\tilde{p}, A)$ must also be assigned positive probability, with $\tilde{p}$ satisfying

$$N(\tilde{p} - \gamma)\psi_\ell - \alpha = N(p_H - \gamma).$$

That is, advertising must be chosen to make sure $\wp$, for type ℓ, is indifferent between $(\tilde{p}, A)$ and (p_H, ϕ). Further, $\tilde{p}$ must be chosen such that only those C for whom $c = g$ purchase, to ensure that if type ℓ is indifferent between (p_H, ϕ) and $(\tilde{p}, A)$, type h strictly prefers $(\tilde{p}, A)$. Finally for type ℓ must must set $(\tilde{p}, A)$ with probability less than 1 so that C is willing to pay a price above p_H. By R5 ≥ 0, however, no such $\tilde{p}$ can be found that also generates a purchase by any C. Therefore this signalling deviation is not feasible either.

Altogether, this argument rules out all possible types of signalling strategies that might destroy the credibility of the updating rule supporting HP. The 9 other cases operate similarly.

Finally, a few words on the extra constraint used to support the AD outcome are required. If one of $N(p_H - \gamma)\psi_h$ or $N(p_L - \gamma)$ exceeds φ's payoff for type h in the AD outcome--$N(\hat{p} - \gamma)\psi_h - \alpha$ --φ may do strictly better (for all types) by adopting a signalling strategy specifying one of (p_H,ϕ) or (p_L,ϕ) with certainty. This would contradict the claimed credibility of the update rule supporting AD as an outcome. The additional restrictions rule out this deviation. No extra restrictions are needed to defeat randomized signalling deviations.

To summarize what has been obtained so far, period 1 outcomes depend on the information consumers have obtained in period 0. If consumers are certain about quality, then, in equilibrium, there is no advertising in period 1 and price is the willingness to pay μ_q. Otherwise, there are three possible outcomes that may be part of a MPSEC: i) HP, in which there is no advertising, and irrespective of quality, price equals the willingness to pay for consumers who had good experience; ii) LP, similar to HP except that price equals the willingness to pay of consumers who had poor experience; or iii) AD, in which φ always advertises and charges a high price when the good is of high quality and, otherwise, randomizes between this same price-ad pair and a low price-no ad pair.

The remaining information needed to study period 0 is φ's period 1 payoff in each of the possible outcomes. If period 0 actions are completely informative, φ earns $N(\mu_q - \gamma)$ given quality q. Otherwise, φ obtains payoffs as follows:

$$\begin{array}{lll} \text{HP:} & q = h & N(p_H - \gamma)\psi_h \\ & q = \ell & N(p_H - \gamma)\psi_\ell \\ \text{LP:} & \text{all } q, & N(p_L - \gamma) \\ \text{AD:} & q = h & N(\hat{p} - \gamma)\psi_h - \alpha \\ & q = \ell & N(\mu_\ell - \gamma) \end{array}$$

with $\hat{p}$ solving $N(\hat{p} - \gamma)\psi_\ell - \alpha = N(\mu_\ell - \gamma)$

2. Period 0

There are two kinds of t=0 outcomes that arise from a MPSEC--one in which advertising occurs with positive probability and one in which it does not. Although much more complicated, the argument used to establish this result is much like that used to analyze period 1.

First, it is shown that if advertising does occur with positive probability in period 0, φ's period 0 strategy must involve type ℓ randomizing over $((p_0, a_0)=)\ (\mu_\ell, \phi)$ and $(\mu_\ell + \frac{\alpha}{N}, A)$, while type h chooses $(\mu_\ell + \frac{\alpha}{N}, A)$ with certainty. Moreover, the updating rule must be such that the observation of $(\mu_\ell + \frac{\alpha}{N}, A)$--which does not permit C to infer q with certainty at the end of period 0--is followed by AD in period 1. That is, a MPSEC in which advertising may occur in period 0 must be followed by AD. (The MPSEC is "AD-AD".)

Next, it is shown that if advertising does not occur with positive probability in period 0 of a MPSEC, φ's period 0 strategy must specify the same (nonrandom) action for both types. This strategy, by itself, is "uninformative" (U)--$\rho(p_0,a_0) = \delta$ -- and so the period 1 outcome is one consistent with the configuration of restrictions given in Table 1. It is then shown that an updating rule yielding HP as the period 1 outcome, following this "uninformative" period 0 behavior, cannot be part of an equilibrium of the game. Thus, if advertising does not occur in period 0, the MPSEC must be of the form U-LP or U-AD, denoting "uninformative behavior followed by LP" and "uninformative followed by AD".

Finally, it is shown that AD-AD, U-LP and U-M are each MPSEC for some parameter values.

It should be noted that the ad game may possess no MPSEC for certain parameter values. Indeed, nonexistence is an issue for PSE in general (see Grossman and Perry). As an example, suppose that U in period 0 yields case 1 in period 1 and the additional restriction needed to support AD in period 1 is not satisfied when $\rho(p_0,a_0) = \delta$. Then the U outcome in period 0 cannot occur. If, for the same parameter values, no period 0 randomization yielding AD at t=1 could generate a $\rho(p_0,a_0)$ for which the restrictions needed to support AD are satisfied, the AD-AD configuration is not a MPSEC either. Since these two types are all the possibilities for MPSEC, none exist for the assumed parameter values. Examples have been produced which generate nonexistence. At the same time, examples have been produced which demonstrate that both kinds of equilibria exist.

Of the several types of potential equilibria, the AD-AD equilibrium is of most interest for the current work. Consequently, the other equilibria are dealt with only briefly and at the end of this section. The analysis of the AD-AD equilibrium begins with a consideration of the features any MPSEC must have should $(p_0,a_0{=}A)$ occur with positive probability.

The first feature of any such equilibrium is that it cannot be that φ always advertises; that is $F_0(p_0,\phi|q) = 0$ for all q and p_0 cannot arise. To see why, consider a period 0 signalling strategy $G(p_0,\phi|q) = F_0(p_0,A|q)$. Here φ behaves exactly as under F_0 except $a_0 = \phi$ replaces $a_0 = A$. Under both F_0 and G any information revelation by φ is accomplished through variations in p_0. Thus C's information upon entering period 1, $\rho_c(p_0,a_0,c)$, is the same irrespective of whether F_0 or G is used in period 1. Under the definition of MPSEC then, it must be the case that F_0 and G are followed by the same second period outcome. Therefore, whether G is a profitable signalling strategy depends on its first period effect only. This effect clearly is to generate lower costs for each p_0. Thus G is a profitable deviation for both types and any updating rule supporting φ always advertising could not be credible.

Second, φ cannot always advertise for one type and never for the other. To support this outcome, φ's payoff must be identical across types. Since type ℓ is revealed by φ's actions, however, φ's payoff for type ℓ (and hence h) cannot exceed $2N(\mu_\ell - \gamma)$ (i.e. $N(\mu_\ell - \gamma)$ in each period). Under any updating rule, both types can improve via a signalling strategy implying $(p_0,a_0) = (\bar{p},\phi)$ with certainty, where $\bar{p} \equiv \delta\mu_h + (1-\delta)\mu_\ell$ (i.e. C's t=0 willingness to pay given only prior information).

Before proceeding further, it should be noted that the situation just excluded is that of separation, the conjectured outcome in Nelson's original model. In that model it is not possible for φ of type ℓ to achieve the same payoff as φ of type h. This occurs because consumption experience is assumed to reveal quality perfectly. In the present model, since quality is not revealed by consumption experience, there is greater scope for imitation of the actions of producers of high quality goods by producers of low quality goods.

Next, whenever φ chooses F_0 such that each type sets one price, p_0^q (which could vary with type), advertising cannot occur with positive probability for either type. The argument establishing this result is quite long due to the fact that even if there is no randomization over prices, there are still a number of signalling strategies to be considered, depending on both the type of randomization over ads hypothesized and the period 1 consequences. The key to the argument, however, is easy. Since φ neither always advertises nor always advertises for one type only, any advertising must be part of a randomized strategy. Further, the strategy must be one in which both types choose the same p_0. The proof shows that for any such strategy that could be part of a MPSEC, there is a signalling strategy involving randomization over both price and ad that makes φ at least as well off for type ℓ and better off for type h.

Taking these preliminary results together, if advertising in period 0 is to be part of MPSEC, φ must randomize over both price and advertising for some type; i.e., there is no MPSEC involving advertising in which φ fails to act in this fashion. The form that this randomization takes must still be determined.

One characteristic that is immediate is that the equilibrium strategy must have some (p_0,a_0) assigned positive probability for both types. If not, then the observed (p_0,a_0) always identifies quality. Given any beliefs consistent with this hypothesized equilibrium, φ could improve by imitating type h when $q = \ell$, in fact, holds.

Now suppose φ's randomization includes a pair (p_0^h,a_0^h) assigned positive probability for $q = h$ but zero probability for $q = \ell$. Let (p_0^ℓ,a_0^ℓ) be chosen by both types, as required by remarks in the previous paragraph. To support the randomization, $\wp$'s payoff for type h must be the same irrespective of whether (p_0^h, a_0^h) or (p_0^ℓ, a_0^ℓ) occurs. Type ℓ's payoff when (p_0^ℓ, a_0^ℓ) arises depends, among other things, on the equilibrium outcome in t=1. If (p_0^ℓ, a_0^ℓ) is followed by HP or AD, type ℓ's payoff falls short of type h's given $\psi_h > \psi_\ell$. However, type h's payoff is available to type ℓ via type ℓ simply choosing (p_0^h, a_0^h); therefore, the hypothesized strategies are not equilibrium strategies. If (p_0^ℓ, a_0^ℓ) is followed by LP, both h and ℓ earn the same payoff. This common payoff will be equal to $2N(\mu_\ell - \gamma)$ should $\wp$ choose some (p_0, a_0) for type ℓ not chosen for type h. In this case the payoff can be improved upon for both types (under any updating rule) by $\wp$ choosing $(\bar{p}, \phi)$ with certainty. If there is no (p_0, a_0) chosen by type ℓ but not type h, then there must be some (p_0, q_0) (possibly (p_0^ℓ, a_0^ℓ) itself) chosen with positive probability by both types that signals $q = \ell$. Then, $\wp$ can do better using a signalling strategy that raises all prices slightly, [i.e. (p_0^h, a_0^h) becomes $(p_0^h + \epsilon(p_0^h), a_0^h)$], leaves advertising behavior unchanged and alters the randomization (again slightly) so that the highest price charged by both types signals $q = \ell$ less strongly. By the assumed continuity of the updating rule in $\rho_c(\cdot)$, this signalling strategy results in LP again in period 1. However, since ρ_c is increased, the t=1 equilibrium price, p_L, is increased. Since this strategy also generates larger expected profits in period 0 for both types, it dominates the proposed equilibrium strategy. Therefore, the beliefs that support such a strategy as an equilibrium strategy cannot be credible.

The above result implies, then, that the (p_0, a_0) pairs possibly chosen if $q = h$ must be a subset of those chosen for $q = \ell$ in any equilibrium. Also, since randomization over both prices and advertising must occur for at least one type, the result just demonstrated implies such randomization occurs for type ℓ at least. This leaves only four possible equilibrium configurations in t=0, distinguished one from the other by φ's strategy given type h; that is, it can either be that $F_0(p_0,a_0|h)$ involves no randomization, randomization over both price and advertising, randomization over price only or randomization over advertising only.

The case in which φ randomizes over both elements of (p_0,a_0) where q=h is easily dealt with. Because φ must randomize over the same set of prices when $q = \ell$, this randomization can possibly be part of an equilibrium only if φ's payoff is independent of q. In particular, it must be that all (p_0,a_0) chosen with positive probability by both types must be followed by LP. However, a signalling strategy involving all prices being raised slightly, similar to that given above, destroys the credibility of the updating rule used to support such a strategy as an equilibrium. This same argument serves to eliminate the situation in which, for type h, φ randomizes over price alone.

When φ randomizes only over advertising, for type h, there must be some (p_0,a_0) that reveals type ℓ. This is a consequence of the fact that randomization must occur over both prices and advertising for some type. Thus φ's payoff for type ℓ is at most $2N(\mu_\ell - \gamma)$. Further, if φ is to randomize over some set of (p_0,a_0) for both types, as occurs here, all such (p_0,a_0) must be followed by LP, since this is the only circumstance in which payoffs are equal across types and (p_0,a_0) pairs. Then, however, φ's payoff for type h also must not exceed $2N(\mu_\ell - \gamma)$, which means that a signalling deviation, G, involving φ choosing $(\overline{p},\phi)$ deterministically for both types yields φ more period 0 profit for every q, and no less in period 1. Thus, φ's randomizing over advertising alone, when q = h, cannot be supported by credible updating either.

Finally, the only remaining possibility is that φ randomizes only if q = ℓ. Could φ's strategy for q = h in this case involve not advertising? Since advertising must then occur for q = ℓ, and q = ℓ is thereby revealed, φ's payoff for type ℓ is at most $N(\mu_\ell - \gamma) - \alpha + (\mu_\ell - \gamma)$. This payoff can be improved upon for type ℓ by choosing (μ_ℓ,ϕ) instead, under any updating rule.

Therefore, all that is left is a period 0 strategy F_0 of the general form initially claimed: φ always chooses a particular $(\tilde{p}_0,\tilde{a}_0)$ if q = h, and randomizes over both prices and advertising for q = ℓ, this randomization including $(\tilde{p}_0,\tilde{a}_0)$. It is immediate that φ's randomization assigns positive probability to exactly one other pair--(μ_ℓ,ϕ)--simply because the price can be no higher and a_0 = A is costly. It follows that $\tilde{a}_0$ = A, for type ℓ to be willing to randomize.

The only remaining task is to show that if the period 0 outcome is to be AD, as described above, AD must be followed by AD in period 1. To do this, suppose first, that AD were followed by LP. φ's payoff for type ℓ is $2N(\mu_\ell - \gamma)$, since observation of (μ_ℓ,ϕ) reveals q = ℓ. As for type h, the period 0 payoff for both types h and ℓ is the same, given $(p_0,a_0) = (\tilde{p}_0,A)$. Under LP, both also earn the same period 1 profit. Thus, if AD is followed by LP and type ℓ is willing to randomize in t=0, φ must obtain a payoff of $2N(\mu_\ell - \gamma)$ for q = h also. This is dominated by a signalling strategy in which $(\bar{p},\phi)$ occurs with certainty for all q.

Suppose AD is followed by HP, and consider the following signalling strategy. For q = h, $\wp$ chooses $(\tilde{p} - \varepsilon, A)$ for ε small; for q = ℓ, $\wp$ randomizes between (μ_ℓ,ϕ) and $(\tilde{p} - \varepsilon, A)$, raising the probability with which (μ_ℓ,ϕ) is observed. By the continuity restriction on beliefs, this signalling strategy is also followed by HP when $(\tilde{p}-\varepsilon,A)$ is observed and, since ρ_C is higher, p_H is increased relative to the proposed equilibrium. It is then easy to check that, should ε be chosen such that $\wp$ is no worse off for type ℓ, $\wp$ is strictly better off for type h.[9] Consequently, the signalling strategy succeeds in showing that the updating rule supporting AD-HP is not credible.

This final result provides one further useful piece of information; namely, that the pairs over which $\wp$ randomizes for q = ℓ must be (μ_ℓ,ϕ) and $(\mu_\ell + \alpha/N, A)$. This outcome occurs because, for type ℓ, $\wp$ must be indifferent between choosing (μ_ℓ,ϕ) and, thereby, revealing q = ℓ, and imitating type h by choosing $(\tilde{p},A)$. The former strategy yields payoff

$$2N(\mu_\ell - \gamma),$$

while the latter gives

$$N(\tilde{p} - \gamma) - \alpha + N(\mu_\ell - \gamma)$$

Equating payoffs gives $\tilde{p} = \mu_\ell + \alpha/N$.

While the analysis so far has shown that, if there is an MPSEC with advertising in t=0, it must be of the form AD-AD, it has not addressed the question of whether (or under what circumstances) AD-AD is, in fact, an equilibrium. Before answering this question it will prove helpful to collect some of the results that have been obtained so far. First, the AD-AD strategy for φ is given by (tilde indicating that this strategy is an equilibrium strategy)

$$\tilde{F}_0(p_0,a_0|h) = \begin{cases} 1 & p_0 \geq \mu_\ell + \alpha/N,\ a_0 = A \\ 0 & \text{otherwise} \end{cases}$$

$$\tilde{F}_0(p_0,a_0|\ell) = \begin{cases} 1 & p_0 \geq \mu_\ell + \alpha/N,\ a_0 = A \\ \tilde{f}_0 & \mu_\ell \leq p_0 < \mu_\ell + \alpha/N,\ a_0 = \phi \\ 0 & \text{otherwise} \end{cases}$$

$$\tilde{F}_1(p_1,a_1|h) = \begin{cases} 1 & p_1 \geq \hat{p},\ a_1 = A \\ 0 & \text{otherwise} \end{cases}$$

and

$$\tilde{F}_1(p_1,a_1|\ell) = \begin{cases} 1 & p_1 \geq \hat{p},\ a_1 = A \\ \tilde{f}_1 & \mu_\ell \leq p_1 < \hat{p},\ a_1 = \phi \\ 0 & \text{otherwise.} \end{cases}$$

Also, the $\tilde{f}_t$ are not unrestricted. Indeed, having observed $(p_0,a_0) = (\mu_\ell + \alpha/N,A)$, C must be willing to purchase, requiring expected utility given observation of $(\mu_\ell + \alpha/N,A)$ be equal to or exceed $\mu_\ell + \alpha/N$:

$$\frac{\delta}{\delta + (1-\tilde{f}_0)(1-\delta)}\mu_h + \frac{(1-\tilde{f}_0)(1-\delta)}{\delta + (1-\tilde{f}_0)(1-\delta)}\mu_\ell \geq \mu_\ell + \alpha/N$$

Similarly, $\tilde{f}_1$ must be such that, having observed (p_0, a_0, p_1, a_1, c) $= (\mu_\ell + \alpha/N, A, \hat{p}, A, g)$, C is willing to purchase:

$$\xi\mu_h + (1-\xi)\mu_\ell \geq \hat{p}$$

with

$$\xi \equiv \frac{\rho_c(\mu_\ell + \alpha, A, g)}{\rho_c(\mu_\ell + \frac{\alpha}{N}, A, g) + (1 - \tilde{f}_1)[1 - \bar{\rho}(\mu_\ell + \frac{\alpha}{N}, A, g)]} ;$$

Finally, for AD to be a period 1 outcome it must be that the HP and LP prices computed using $\rho(\mu_\ell + \alpha/N, A)$ (which depends on $\tilde{F}_0$) satisfy

$$N(\hat{p} - \gamma)\psi_h - \alpha > N \max\{(p_H - \gamma)\psi_h,\ p_L - \gamma\} .$$

Now, can the AD-AD strategy $\tilde{F}_t$ be supported as a MPSEC? The answer depends on joint restrictions on parameter values and the updating rule. Suppose, for instance, that it is assumed that, should t=0 actions be uninformative (so that $\rho(p_o, a_o) = \delta$), LP is a possible period 1 outcome. Also, assume that the following condition holds:

$$(\mu_\ell - \gamma)(\frac{\psi_h}{\psi_\ell} + 1) + \frac{\alpha}{N}(\frac{\psi_h}{\psi_\ell} - 1) \geq \bar{p} - \gamma + p_L - \gamma$$

This inequality is equivalent to the requirement that the payoff to φ under the strategy AD-AD, when q=h, is larger than the payoff that φ could obtain from a signalling strategy equivalent to the F_t making up the U-LP outcome (with p_L computed using $\rho(p_0, a_0) = \delta$).

The updating rule is already restricted by the continuity assumption. In the current circumstances, this assumption requires that, for any (p_0, a_0) yielding a value of $\rho(p_0, a_0)$ sufficiently close to the value implied by $\tilde{F}_0$, C's beliefs in period 1, obtained by credible updating given (p_0, a_0, c), must generate AD as the period 1 outcome. Suppose, in addition, the updating rule specifies that for other (p_0, a_0), beliefs supporting LP are used if possible. When LP is not a possible period 1 outcome, beliefs should support HP if possible, and AD otherwise. Beliefs in t=0 are allowed to be any that support AD-AD as a sequential equilibrium.

Given the above parameter restriction, it can be shown that $\tilde{F}_t$ coupled with any updating rule satisfying the restriction just given, constitute a MPSEC. The complete argument demonstrating this proposition is quite long. The basic elements, however, are fairly easy. In particular, it is plain that there are a variety of beliefs that, when coupled with $\tilde{F}_t$, form a SE. The issue, therefore, is simply one of the credibility of the updating rule. The proof that AD-AD is a MPSEC must show that there are no signalling strategies, G, that destroy credibility of the proposed updating rule. The point can be demonstrated as follows. Consider pure signalling strategies first. As above, it is easy to check that pure strategies that separate types are not sustainable; φ will always deviate for $q = \ell$. Thus, both types use the same pure strategy. Since C's information is not altered by such a deviation (i.e. $\rho(p_0, a_0) = \delta$), the restriction that the updating rule specifies that

U be followed by LP applies. For type ℓ, φ earns a greater payoff from such a deviation. For type h, however, the assumed parameter restriction implies a reduced payoff. Thus no pure signalling strategy can undermine credibility.

Randomized signalling strategies also must be considered. A long argument, very much like that used to show that AD-AD is the only possible MPSEC involving advertising, establishes that there is no randomized signalling strategy that will destroy the credibility of the updating rule. Put simply, $\tilde{F}_t$ is the most profitable randomized strategy for φ, if q = h, that is also sustainable in the sense that type ℓ is willing to randomize in the manner in which signalling strategy requires.

Two other points concerning the AD-AD equilibrium should be mentioned at this point. First, other parameter restrictions and update rules will also yield AD-AD as an equilibrium. For example, suppose the update rule specifies (for t=1) HP when possible, then LP failing that, then AD. Assume also that parameters are such that i) HP is a possible period 1 outcome given $\rho(p_0,a_0) = \delta$ and ii) φ's payoff, given q = h, from AD-AD dominates U-HP. Then $\tilde{F}_t$ is again a MPSEC.

Second, if for uninformative period 0 outcomes, AD exists and is the only possibility in period 1, $\tilde{F}_t$ can never be MPSEC. The argument is easy. A signalling strategy involving uninformative period 0 behavior yields the same period 1 profits as AD-AD for both q. Thus whether the signalling strategy pays depends only on period 0 payoffs. Under AD-AD, both types earn a first period payoff valued at $N(\mu_{\ell} - \gamma)$, which falls short of the $N(\bar{p} - \gamma)$ achievable by the above signalling strategy. Thus, for AD-AD to be a MPSEC, it must be possible for some outcome other than AD to occur in period 1.

The intuition for this last result is easy. The point of costly advertising in period 0 is to ensure greater profits later, by advertising yet again. However, if those period 1 profits are necessarily obtained (i.e. AD must be the period 1 outcome) following some more profitable period 0 action, there is no return to sacrificing current payoff.

As mentioned at the beginning of this section, there is another type of MPSEC, in which $a_0 = A$ does not occur with positive probability. The fact that advertising does not occur at t=0 is an indication that, should φ be of type h, he does not find attempts to distinguish himself from type ℓ profitable enough to compensate for the ad cost. As might be expected in such a situation, the equilibrium outcome that results is a pure pooling equilibrium. The most profitable of these equilibria for both types is obviously one yielding $(\bar{p},\phi)$ at t=0. The remainder of this section considers equilibria of this type.

One set of equilibrium strategies associated with a t=0 pooling outcome is as follows:

$$\tilde{F}_0(p_0,a_0|q) = \begin{cases} 1 & p_0 \geq \bar{p},\ a_0 = \phi, \text{ all } q \\ 0 & \text{otherwise,} \end{cases}$$

and

$$\tilde{F}_1(p_1,a_1|q) = \begin{cases} 1 & p_1 \geq \bar{p}_L,\ a = \phi, \text{ all } q \\ 0 & \text{otherwise.} \end{cases}$$

The updating rule may be any rule that yields LP whenever possible in period 1 and either HP or AD otherwise.

As usual, it is trivial to show that such $\tilde{F}_t$ are part of a SE, so the issue is again credibility of updating. As in the demonstration that AD-AD is a MPSEC, any sustainable pure signalling strategy must involve φ charging the same price for both types; therefore, whether the deviation is more profitable than the equilibrium depends only on the period 0 payoff. However, $(\bar{p},\phi)$ is the most profitable period 0 action. Thus a successful signalling strategy must alter *C*'s information via randomization.

Because φ's payoff for type ℓ dominates the payoff that would be earned if φ revealed type ℓ, the signalling deviation must not reveal type ℓ. Also, it is easy to check that a signalling strategy in which φ reveals type h is not sustainable. Taking these facts together, then, either both types randomize over the same actions or else the signalling strategy involves type ℓ randomizing over the equilibrium action $(\bar{p},\phi)$ and some other (p_o^*,a_o^*) which type h chooses with certainty.[10] The first option must contain some outcome signalling that type ℓ is the more likely type. This reduces φ's payoffs for type ℓ (at least) in both periods because p_1 is either a lower value of p_L or p_H or the price is μ_ℓ while p_0 must also be lower to induce *C* to purchase. In the second outcome (p_0^*,a_0^*) indicates q = h. Referring to Table 1, if LP may follow, for $\rho(p_o,a_o) = \delta$, it can also follow any (p_0,a_0) for which $\rho(p_0,a_0) > \delta$ (as is true of (p_0^*,a_0^*)). Thus, by the assumed updating rule (p_0^*, a_0^*) is followed by LP, which means that φ receives the same payoff from the signalling strategy irrespective of type. If this payoff exceeds that earned by following the equilibrium path, type ℓ will not wish to randomize between $(\overline{p},\phi)$ and (p_0^*, a_0^*); i.e. the signalling strategy is not sustainable. Otherwise, the deviation fails to improve (strictly) for some type. In either case, the signalling strategy fails to destroy credibility.

An argument virtually identical to that just given shows that it is never possible to construct MPSEC in which U-HP occurs. It is always possible to construct a randomized signalling strategy much like the one just described, and in which φ is made no better off for type ℓ. This deviation yields $\rho(p_0^*, a_0^*) = \delta + \varepsilon$, for ε small. Since type h's payoff exceeds type ℓ's in the HP outcome, the signalling strategy improves strictly for type h, thereby destroying credibility.

Whether U-AD can be part of a MPSEC is more complex and depends on parameter values. Again, a signalling strategy must involve randomization. Moreover, AD cannot be the second period outcome following any (p_0^*, a_0^*) allowed by the signalling strategy because U-AD is the most profitable of such situations for type ℓ. Similarly, if some (p_0^*, a_0^*) were followed by LP, then any circumstance in which φ is willing to randomize for type ℓ can be shown to make type h worse off relative to the equilibrium strategy. The only possibility for overturning the credibility of the updating rule, therefore, is a signalling strategy, resulting in HP and in which type ℓ is willing to randomize between $(\overline{p},\phi)$ and some (p_0^*,a_0^*), the latter point being such that for type h, φ strictly prefers choosing it to $(\overline{p},\phi)$. Examples can be constructed in which this possibility may occur, as well as cases for which it may not.

IV. PREDICTIONS

What can be learned from this model of advertising? Observable implications come in three kinds.

The first involves basic features of the model's equilibrium. As indicated above, a major empirical anomaly in the context of the earlier literature is the existence of recurrent advertising. The model developed above not only explains the existence of repeat advertising but, in fact, requires that it be part of the subsequent strategic behavior of any firm that advertises when the good is introduced.

In addition to explaining the anomalous repeat advertising, the model also generates a simple and robust new prediction: whenever advertising is observed after a good is introduced (i.e. advertising occurs in period 1), the price at which the good is sold must exceed that at which it was sold when introduced. That quantity sold declines is also implied.

This prediction holds for both equilibria that permit $a_1 = A$: U-AD and AD-AD. The proof is simply that in the U-AD equilibrium $p_0 = \overline{p} < p_H < \hat{p}$, where the first inequality holds because c = g is "good news", and the second is required to prevent a signalling deviation in which φ always charges p_H from undermining the AD outcome in period 1. In the AD-AD equilibrium, p_0 is at most $\mu_\ell + \alpha/N$ and $\mu_\ell + \alpha/N < \gamma + (\mu_\ell + \alpha/N-\gamma)/\psi_\ell = \hat{p}$, the period 1 price when $a_1 = A$. The exact intuition behind this result depends on the equilibrium. In U-AD, $p_1 > p_0$ because period 1 is the earliest time at which φ signals quality, with $a_1 = A$ indicating high quality and so permitting φ to change a higher price. In AD-AD, while φ may also signal quality in period 0, in period 1, $a_1 = A$ yields fewer customers (those who observed c = b no longer buy). Therefore p_1 must be higher to yield the equality of period 1 payoffs across advertising outcomes required to support φ's randomization for $q = \ell$.

Another simple prediction is that when $a_0 = A$ is observed (thus ruling out U-AD as the equilibrium generating the data), period 1 profits must not fall short of period 0 profits, provided $a_1 = A$ is observed. That is, repeated advertising cannot be accompanied by declining (per period) profits. The proof is simply that in the AD-AD equilibrium, φ earns profits equal to $N(\mu_\ell - \gamma)$, in period 0. When $a_1 = A$, profits are equivalent to $N(\mu_\ell - \gamma)$, for type ℓ, and $N(\mu_\ell - \gamma) + (\psi_h - \psi_\ell)(\hat{p} - \gamma) > N(\mu_\ell - \gamma)$, for type h, are earned in period 1. Thus, profits do not fall, and may rise, over time if repeat advertising is observed.

These predictions are particularly useful because they involve few variables and are cast in a manner that renders testing comparatively straightforward. That is, they refer to time series observations on a given firm, in which case many difficult-to-measure entities $(\delta,\gamma,\ldots)$ may be ignored so long as they may plausibly be assumed constant over time. Moreover, the predictions are sufficiently elementary that it may reasonably be expected that they are robust to a variety of relaxations of the model's restrictions.

The second type of prediction involves cross-section comparisons of firms. These implications are somewhat more limited because very little can be said unless there is some mechanism that determines which equilibrium any given firm and group of consumers are playing. Of course, assumptions can be made that allow derivation of some implications--for example, it might be supposed that the equilibrium is picked by a coin toss or is at least not systematically related to quality or other parameters of the model. However, the conclusions obtained in this fashion are heavily dependent on the assumed procedure, which is essentially arbitrary.

A more promising route appears to be to once again make use of the model's structure and only examine firms for which $a_0 = A$. It then follows once again that AD-AD is the equilibrium generating the data. A stochastic variant of the basic Nelson proposition then emerges. Advertising expenditures $\alpha(a_1)$ are positively correlated with product quality (the correlation being across games having identical parameters, $\delta, \gamma, \psi_h, \psi_\ell, V$ and α). Moreover, so are product price (p_1) and period 1 profits π_1. The quantity sold in period 1 is negatively correlated with $\alpha(a_1)$. Note that a variety of specific values of these covariances are consistent with the model because there are many possible equilibrium randomizations. All of course have the same fundamental structure. Thus, for given parameters, only bounds on the magnitudes of the covariances are implied by the theory. Checking whether the data are consistent with those bounds requires measurement of the underlying parameters, perhaps a difficult process. Nevertheless, the sign restrictions can be checked more easily because they do not depend on the underlying parameter values. That is, the signs of the unconditional covariances are the same as those of the conditional ones. Thus sign restrictions, at least, are arguably checkable.

If data on the underlying parameters are available, the third type of proposition, comparative statics, may also be approached. Again, provided conditioning on $a_0 = A$ is imposed, so that AD-AD is the equilibrium, some clear results emerge. The endogenous variables are p_1, period 1 profits π_1, and quantity sold in period 1, Q.

Consider a change in the technology of home consumption such that the utility of c = g, c = b, or both, rises. μ_q thus increases, and since $p_1 = \hat{p}$ or μ_ℓ, p_1 must rise. Similarly, π_1 increases and Q is left unchanged (at $\psi_q N$ or N).

More interesting is the effect of an increase in marginal cost γ. Since $\partial\hat{p}/\partial\gamma = (1 - 1/\psi_\ell) < 0$, p_1 either falls (if $\hat{p} = p_1$) or remains unchanged (if $p_1 = \mu_\ell$) as γ increases: a somewhat unusual outcome. π_1 does the same, and again, Q is unchanged. $\hat{p}$ must fall because $\hat{p}$ is the price at which ϕ of type ℓ is indifferent between selling to all N consumers at price $p_1 = \mu_\ell$ and selling to the $\psi_\ell N$ consumers who had good period 0 consumption experience. In the latter case, when γ rises, the cost increase applies to fewer consumers, in which case the price obtained must fall to achieve indifference for ϕ; i.e. the payoffs to the $(\hat{p},A)$ and (μ_ℓ,ϕ) outcomes in period 1 must fall equally.

Also, note that period 1 outcomes do not depend on δ except insofar as δ affects whether AD-AD is the equilibrium. Period 1 price, profit and quantity sold are entirely determined by the requirement that ϕ is willing to randomize when $q = \ell$. Thus, a strong "zero restriction" is that all else constant, p_1, π_1 and Q are independent of δ if $a_1 = A$.

The influence of changes in ψ_h is straightforward. Since ψ_h is not of interest to φ when $q = \ell$, p_1 does not depend on ψ_h, irrespective of a_1. However, since greater ψ_h increases φ's sales when $a_1 = A$ and $q = h$, on average π_1 will rise with ψ_h given $a_1 = A$, as will Q.

On the other hand, increments to ψ_ℓ raise p_1 when $a_1 = \phi$ $(p_1 = \mu_\ell)$ and may either augment or lower p_1 for $a_1 = A$ $(p_1 = \hat{p})$:

$$\frac{\partial \hat{p}}{\partial \psi_\ell} \propto (\gamma - V - \alpha/N),$$

which may take on either sign. Since $\pi_1 = N(\mu_\ell - \gamma)$ for type ℓ and $\mu_\ell = 1 + (1-V)\psi_\ell$, φ is always made better off by greater ψ_ℓ when $q = \ell$. But for $q = h$, whether φ receives a greater payoff depends only on the sign of the change in $\hat{p}$. Thus, for example, in contrast to what might be expected, it may not be to φ's advantage to have ψ_ℓ low when $q = h$ in fact holds.

Finally, consider an increment to the cost of advertising, α. p_1 must rise for $a_1 = A$, so as to yield φ the same period 1 payoff as does not advertising, for $q = \ell$. For $a_1 = A$, $p_1 = \mu_\ell$ holds, so p_1 is independent of α in that case. For $q = h$,

$$\frac{\partial \pi_1}{\partial \alpha} = \frac{\psi_h}{\psi_\ell} - 1 > 0,$$

in which case producers of high quality goods are made better off by an increment to ad costs. The reason is simply that if α is increased, a price adjustment must occur for advertising to continue serving its informational role. This adjustment is required for type ℓ to be willing to advertise and still be no worse off than he would be were quality simply (and truthfully) announced. However, the price adjustment is large enough to yield ϕ a strictly higher payoff for $q = h$. Last, Q is unaffected by the change in α.

V. CONCLUSIONS

In this model, the issue investigated was whether repeat advertising should be observed in equilibrium. The focus on the time aspects of advertising is really not crucial, however. The important point is that consumers have imperfect information regarding the quality of the period one good and that consumption of the period zero good provides information about the period one good's quality (as do period zero ads). Therefore, the issue addressed by the model might equally well be one regarding the advertising of two different products and the value of brand name advertising should both products be produced by the same firm. This problem is a subject of current research.

Also, the reader should note that this model provides a formalization of the old notion of advertising as an investment in "goodwill". Consumers in this model acquire informational capital on quality and advertising serves as a means of increasing the consumers' capital stock. As such, ads enable the firm to set a higher period one price than would be possible if no advertising occurred. As well, the informational capital of consumers provides an explanation for possible differences between incumbents and new entrants, making possible an equilibrium treatment of the issue of advertising and entry.

FOOTNOTES

[1]Batteries are a good example. Checking durability before purchase is difficult, as is inferring it from a single dead battery.

[2]This point is the focus of Kihlstrom and Riordan (1984) and Milgrom and Roberts. Producers of high quality goods advertise, while producers of low quality goods do not. This behavior makes advertising (along with price, in Milgrom and Roberts) a perfect signal of quality.

[3]To illustrate, why does Pepsi hire Michael Jackson instead of simply setting a pile of cash ablaze? It is obvious that Michael Jackson has excellent alternatives and so must be paid a large sum to perform for Pepsi. A stock of cash may readily be faked and so is not obviously a costly activity.

[4]Recall that the firm knows the value of q while initially the consumer does not.

[5]With some additional notation, production costs can be made positively related to quality.

[6]$B_1(b) \subset B_1(g)$ follows from $\psi_h > \psi_\ell$, and must hold for (p_1, a_1) that occur with positive probability in equilibrium. Here it is required that the inclusion also holds off the equilibrium path, that is, that beliefs are not directly dependent upon consumption.

[7]Here it is assumed that, should $\rho_c = 1$ or 0, no additional price or advertising signals alter this probability. This assumption is without loss of generality given the assumed structure of strategies and the equilibrium concept.

[8]Type ℓ is as well off under the deviation as along the equilibrium path. However, the change in $\wp$'s payoff for $q = h$ is

$$[N\psi_h(\hat{p} - \gamma) - \alpha] - [N\psi_h(\tilde{p} - \gamma)],$$

which simplifies to $\alpha(\psi_h/\psi_\ell - 1) > 0$.

[9]To obtain this result, note that if type ℓ is no worse off, then

$$(\tilde{p} - \gamma)N - \varepsilon N - \alpha + [p_H(\varepsilon) - \gamma]\psi_\ell N = (\tilde{p} - \gamma)N - \alpha + (p_H - \gamma)\psi_\ell N,$$

or $\quad \varepsilon = \psi_\ell[p_H(\varepsilon) - p_H]$.

This observation implies that for $q = h$, profits from the signalling strategy are

$$(\tilde{p} - \gamma)N - \psi_\ell[p_H(\varepsilon) - p_H]N - \alpha + [p_H(\varepsilon) - \gamma]\psi_h N > (\tilde{p} - \gamma)N - \alpha$$
$$+ (p_H - \gamma)\psi_h N$$

since $[p_H(\varepsilon) - p_H]\psi_h - [p_H(\varepsilon) - p_H]\psi_\ell > 0$.

[10]Of course (p_0^*, a_0^*) may be some set of points, each element of which h chooses with positive probability.

REFERENCES

Banks, Jeffrey S., and Sobel, Joel. "Equilibrium Selection in Signalling Games." Econometrica 55 (1987), 647-61.

Cho, In-Koo, and Kreps, David M. "Signalling Games and Stable Equilibria. Quarterly Journal of Economics 102 (May 1987), 179-221.

Farrell, Joseph. "Credible Neologisms in Games of Communication." Mimeo, June 1985.

Grossman, Sanford J., and Perry, Motty. "Perfect Sequential Equilibrium." Journal of Economic Theory 39 (1986), 97-119.

Kihlstrom, Richard E., and Riordan, Michael H. "Advertising as a Signal." Journal of Political Economy 92 (June 1984), 427-50.

Kreps, David M., and Wilson, Robert. "Sequential Equilibria." Econometrica 50 (June 1982), 863-94.

Milgrom, Paul, and Roberts, John. "Price and Advertising Signals of Product Quality." Journal of Political Economy 94 (August 1986), 796-821.

Nelson, Phillip. "Advertising as Information." Journal of Political Economy 82 (July/August 1974), 729-54.

Rochester Center for Economic Research
University of Rochester
Department of Economics
Rochester, NY 14627

1986-87 DISCUSSION PAPERS

WP#46 FORMULA BUDGETING: THE ECONOMICS AND ANALYTICS OF FISCAL POLICY UNDER RULES, by Eric A. Hanushek, June 1986

WP#48 EXCHANGE RATE POLICY, WAGE FORMATION, AND CREDIBILITY by Henrik Horn and Torsten Persson, June 1986

WP#49 MONEY AND BUSINESS CYCLES: COMMENTS ON BERNANKE AND RELATED LITERATURE, by Robert G. King, July 1986

WP#50 NOMINAL SURPRISES, REAL FACTORS AND PROPAGATION MECHANISMS by Robert G. King and Charles I. Plosser, Final Draft: July 1986

WP#51 JOB MOBILITY IN MARKET EQUILIBRIUM by Glenn M. MacDonald, August 1986

WP#52 SECRECY, SPECULATION AND POLICY by Robert G. King, (revised) August 1986

WP#53 THE TULIPMANIA LEGEND by Peter M. Garber, July 1986

WP#54 THE WELFARE THEOREMS AND ECONOMIES WITH LAND AND A FINITE NUMBER OF TRADERS, by Marcus Berliant and Karl Dunz, July 1986

WP#55 NONLABOR SUPPLY RESPONSES TO THE INCOME MAINTENANCE EXPERIMENTS by Eric A. Hanushek, August 1986

WP#56 INDIVISIBLE LABOR, EXPERIENCE AND INTERTEMPORAL ALLOCATIONS by Vittorio U. Grilli and Richard Rogerson, September 1986

WP#57 TIME CONSISTENCY OF FISCAL AND MONETARY POLICY by Mats Persson, Torsten Persson and Lars E. O. Svensson, September 1986

WP#58 ON THE NATURE OF UNEMPLOYMENT IN ECONOMIES WITH EFFICIENT RISK SHARING, by Richard Rogerson and Randall Wright, September 1986

WP#59 INFORMATION PRODUCTION, EVALUATION RISK, AND OPTIMAL CONTRACTS by Monica Hargraves and Paul M. Romer, September 1986

WP#60 RECURSIVE UTILITY AND THE RAMSEY PROBLEM by John H. Boyd III, October 1986

WP#61 WHO LEAVES WHOM IN DURABLE TRADING MATCHES by Kenneth J. McLaughlin, October 1986

WP#62 SYMMETRIES, EQUILIBRIA AND THE VALUE FUNCTION by John H. Boyd III, December 1986

WP#63 A NOTE ON INCOME TAXATION AND THE CORE by Marcus Berliant, December 1986

WP#64 INCREASING RETURNS, SPECIALIZATION, AND EXTERNAL ECONOMIES: GROWTH AS DESCRIBED BY ALLYN YOUNG, By Paul M. Romer, December 1986

WP#65 THE QUIT-LAYOFF DISTINCTION: EMPIRICAL REGULARITIES by Kenneth J. McLaughlin, December 1986

WP#66 FURTHER EVIDENCE ON THE RELATION BETWEEN FISCAL POLICY AND THE TERM STRUCTURE, by Charles I. Plosser, December 1986

WP#67 INVENTORIES AND THE VOLATILITY OF PRODUCTION by James A. Kahn, December 1986

WP#68 RECURSIVE UTILITY AND OPTIMAL CAPITAL ACCUMULATION, I: EXISTENCE, by Robert A. Becker, John H. Boyd III, and Bom Yong Sung, January 1987

WP#69 MONEY AND MARKET INCOMPLETENESS IN OVERLAPPING-GENERATIONS MODELS, by Marianne Baxter, January 1987

WP#70 GROWTH BASED ON INCREASING RETURNS DUE TO SPECIALIZATION by Paul M. Romer, January 1987

WP#71 WHY A STUBBORN CONSERVATIVE WOULD RUN A DEFICIT: POLICY WITH TIME-INCONSISTENT PREFERENCES by Torsten Persson and Lars E.O. Svensson, January 1987

WP#72 ON THE CONTINUUM APPROACH OF SPATIAL AND SOME LOCAL PUBLIC GOODS OR PRODUCT DIFFERENTIATION MODELS by Marcus Berliant and Thijs ten Raa, January 1987

WP#73 THE QUIT-LAYOFF DISTINCTION: GROWTH EFFECTS by Kenneth J. McLaughlin, February 1987

WP#74 SOCIAL SECURITY, LIQUIDITY, AND EARLY RETIREMENT by James A. Kahn, March 1987

WP#75 THE PRODUCT CYCLE HYPOTHESIS AND THE HECKSCHER-OHLIN-SAMUELSON THEORY OF INTERNATIONAL TRADE by Sugata Marjit, April 1987

WP#76 NOTIONS OF EQUAL OPPORTUNITIES by William Thomson, April 1987

WP#77 BARGAINING PROBLEMS WITH UNCERTAIN DISAGREEMENT POINTS by Youngsub Chun and William Thomson, April 1987

WP#78 THE ECONOMICS OF RISING STARS by Glenn M. MacDonald, April 1987

WP#79 STOCHASTIC TRENDS AND ECONOMIC FLUCTUATIONS by Robert King, Charles Plosser, James Stock, and Mark Watson, April 1987

WP#97 THE TRADE-OFF BETWEEN CHILD QUANTITY AND QUALITY: SOME EMPIRICAL EVIDENCE
by Eric Hanushek, September 1987

WP#98 SUPPLY AND EQUILIBRIUM IN AN ECONOMY WITH LAND AND PRODUCTION
by Marcus Berliant and Hou-Wen Jeng, September 1987

WP#99 AXIOMS CONCERNING UNCERTAIN DISAGREEMENT POINTS FOR 2-PERSON BARGAINING PROBLEMS
by Youngsub Chun, September 1987

WP#100 MONEY AND INFLATION IN THE AMERICAN COLONIES: FURTHER EVIDENCE ON THE FAILURE OF THE QUANTITY THEORY
by Bruce Smith, October 1987

WP#101 BANK PANICS, SUSPENSIONS, AND GEOGRAPHY: SOME NOTES ON THE "CONTAGION OF FEAR" IN BANKING
by Bruce Smith, October 1987

WP#102 LEGAL RESTRICTIONS, "SUNSPOTS", AND CYCLES
by Bruce Smith, October 1987

WP#103 THE QUIT-LAYOFF DISTINCTION IN A JOINT WEALTH MAXIMIZING APPROACH TO LABOR TURNOVER
by Kenneth McLaughlin, October 1987

WP#104 ON THE INCONSISTENCY OF THE MLE IN CERTAIN HETEROSKEDASTIC REGRESSION MODELS
by Adrian Pagan and H. Sabau, October 1987

WP#105 RECURRENT ADVERTISING
by Ignatius J. Horstmann and Glenn M. MacDonald, October 1987

To order copies of the above papers complete the attached invoice and return to Christine Massaro, W. Allen Wallis Institute of Political Economy, RCER, 109B Harkness Hall, University of Rochester, Rochester, NY 14627. Three (3) papers per year will be provided free of charge as requested below. Each additional paper will require a $5.00 service fee which must be enclosed with your order. For your convenience an invoice is provided below in order that you may request payment from your institution as necessary. Please make your check payable to the **Rochester Center for Economic Research**. **Checks must be drawn from a U.S. bank and in U.S. dollars**.

W. Allen Wallis Institute for Political Economy

Rochester Center for Economic Research, Working Paper Series

OFFICIAL INVOICE

Requestor's Name ______________________________

Requestor's Address ______________________________

____

____

____

Please send me the following papers free of charge (**Limit: 3 free per year**).

WP# ______ WP# ______ WP# ______

I understand there is a $5.00 fee for each additional paper. Enclosed is my check or money order in the amount of $__________. Please send me the following papers.

WP# ______ WP# ______ WP# ______

WP# ______ WP# ______ WP# ______

WP# ______ WP# ______ WP# ______

WP# ______ WP# ______ WP# ______