

Predation, Reputation, and Entry Deterrence*

PAUL MILGROM

*J. L. Kellogg Graduate School of Management,
Northwestern University, Evanston, Illinois, 60201*

AND

JOHN ROBERTS

Graduate School of Business, Stanford University, Stanford, California 94305

Received July 30, 1980; revised July 13, 1981

Economists often argue that predatory practices are irrational, since there exist cheaper or more certain means to gain or maintain a monopoly. Our game-theoretic, equilibrium analysis suggests that if a firm is threatened by several potential entrants, then predation may be rational against early entrants, even if it is costly when viewed in isolation, because it yields a reputation which deters other entrants. Asymmetric information plays a crucial role in our analysis, since it provides the rationale for entrants to base their expectations of the firm's future behavior on its past actions. The analysis also suggests methods to treat general reputational phenomena. *Journal of Economic Literature* Classification Numbers: 611, 022, 026, 612.

1. INTRODUCTION

Allegations of price cutting or similar tactics aimed at driving a rival out of business are frequently heard, both in and out of law courts. Yet a large fraction of the economics profession would argue that such predation is an irrational strategy for attempting to gain or maintain a monopoly position and that it is, therefore, unlikely to be adopted in practice. This position rests on arguments that predation is costly to the predator and is unlikely to

* This work was begun while Roberts was on the faculty at Northwestern and completed while Milgrom was visiting Stanford. We would like to thank Robert Wilson and David Kreps for making their unpublished results available to us and for several helpful conversations, David Besanko and Garth Saloner for their excellent research assistance, the referees for their useful comments, and the Center for Advanced Study in Managerial Economics, the Office of Naval Research (N00014-79-C-0685) and the National Science Foundation (SOC 79-07542, SES 80-01932, and SES 81-08226) for financial support.

succeed in driving out a rival who understands that the price cutting is temporary. Further, it is held that even if the rival is eliminated, any attempt to raise prices so as to reap the benefits of the monopoly position may attract new entrants. Thus, any monopoly gains would be short-lived.¹ (See, e.g., McGee [10, 11] for full expositions of these arguments, as well as some indication of the nature of the opposing views.)

In this paper we present a model in which predation emerges as a rational, profit-maximizing strategy. In this model, predation is practiced not because it is directly profitable to eliminate the particular rival in question, but rather because it may deter future potential entrants. The mechanism by which this deterrent effect comes about is that by practicing predation the firm establishes a reputation as a predator. This reputation then leads potential entrants to anticipate that the incumbent firm will behave similarly if they should enter, and, thus, entry appears less attractive to them.

In this context, it is worth noting that predation will emerge in our model even if, as asserted by those who doubt the rationality and relevance of predatory strategies, predation against a particular rival involves losses that cannot be directly recouped in the given market, even were exit to be induced. Moreover, viability of this predatory strategy does not depend on being able to induce exit. Rather, all that is needed is that the predator usually be able to drive the rival's return from entry below that available elsewhere.

Examples consistent with the sort of analysis we will develop are not hard to find. Government studies in the U.S. and the U.K. in the early part of the century identified many instances of predatory pricing against new entrants in the ocean shipping industry through the use of "fighting ships," and the U.S. Department of Justice [18] has documented more recent episodes in this industry which it views as predatory (see also Yamey [19]). Although it is difficult to determine if any of this price cutting was done with a view to deterring future entrants, one might expect that firms considering entering an industry with a century-long history of aggressive responses to entry would at least entertain the idea that they might meet a similar response. More directly, Bröck's discussion [2] of IBM's pricing and product strategies against the "plug compatible" manufacturers (who were marketing peripheral equipment for use with IBM central processing units) suggests that IBM was concerned that failure to respond aggressively would encourage further entry. Finally, the fierce price wars that erupted as Proctor and Gamble introduced its Folger's brand of coffee into local markets in the Eastern U.S. in competition with Maxwell House may well have been central in *P & G's*

¹ This latter argument seems to suggest the possibility of deterring entry by holding current prices down. For a discussion of some of the problems with this notion of limit pricing, see [13].

decisions not to continue expansion of its distribution area and, in particular, not to enter the New York City market.²

Of course, it has long been recognized in the literature on industrial organization that the response that entrants expect from incumbent firms would be a major factor in determining the attractiveness of entry, and much of the traditional literature on entry deterrence effectively hinges on the threat of predation (see, e.g., Dixit [3] or Spence [17]). Most models involving such threats are, however, subject to a telling, fundamental criticism. One should expect that the threat of predation will be effective in preventing entry only if entrants find the threat credible. But in these models, which involve a single entrant, if the entrant were to call the incumbent's bluff by entering despite the threat, the incumbent would not be willing to prey, since sharing the market would typically be more the profitable course. In contrast, the strategy of predation in our model does not involve threats which would not rationally be carried out, since the immediate losses incurred in predation are offset by the gains from a reduced threat of further entry that building a predatory reputation yields.

Thus, in our framework, a multiplicity of potential entrants plays a key role in rationalizing a predatory strategy. If, in fact, the established firm were to face a countable infinity of threats of entry, all of which may be potentially carried out, then it is trivial to identify an equilibrium which involves a credible threat of predation, as we demonstrate in Appendix A. (We also show there that there are numerous other equilibria in this framework.) If, however, there are only a finite number of potential entrants, the issue is more complicated. This point has been made by Selten [16].

Selten considers a model which may be interpreted in terms of a firm which operates in N identical markets (a chain store). Each market has one potential entrant. Sequentially, the entrants must individually decide whether to enter the corresponding markets. If entry does not take place in market n (at stage n), the incumbent enjoys its monopoly position in that market without further threat. This contributes Π^M to its overall payoff. If entry occurs at stage n , the incumbent must decide whether to prey on the entrant (yielding Π^P) or to share the market (yielding Π^C). We assume $\Pi^M > \Pi^C > \Pi^P$: predation is costly. In any event, the next entrant must then make its decision, knowing the history of play up through the preceding stage. Assume that the payoff to any entrant from meeting predation is strictly less than that from staying out, which, in turn, is less than the payoff from entry if no predation occurs.

Selten suggests, and it does seem intuitively appealing, that in early rounds of this game the incumbent would adopt the costly predatory action in order

² In Section 4 of their companion paper, Kreps and Wilson present a model which predicts behavior that closely approximates that in the Folger's-Maxwell House competition in several important aspects.

to persuade later entrants that they should best stay out, and that only near the end of the game would it be willing to share a market. Yet Selten also points out that this strategy cannot be consistent with the natural solution concept (that of perfect equilibrium) to employ in such situations.

The argument is the following. Consider the last entrant. It knows that if it enters and meets predation, it would have been better off to stay out. But it also knows that, if faced by actual entry, the established firm is strictly better off if it behaves nonaggressively. Thus, assuming that both firms will always act in their own best interests, entry will occur in the last market to be threatened and will meet a nonaggressive response. Moreover, this will be the result, no matter what has been the history of play to this point.

Now consider the second-last market to be threatened. If entry were to occur there and if the chain store could deter entry in the last market by adopting predatory practices, it might well adopt such measures. However, as just shown, the outcome in the last market is completely determined, independent of the outcome in the second-last market. Thus, if entry occurs, the chain store will share the second-last market peacefully, and, thus, too, entry will occur in this market.

The induction is inexorable and the conclusion clear: in equilibrium, predation will never be practiced. Moreover, even if (for whatever reason) the chain store were observed to have preyed repeatedly against every previous entrant, the logic still will lead the next entrant to anticipate not that past behavior will be repeated but rather that it entry will meet a nonaggressive response. Repeated observations of behavior which, *a priori*, the entrant expected never to see cannot and do not shake its absolute confidence in its predictions of future behavior.

The key factor driving this conclusion is that it is *common knowledge* (see Aumann [1], Milgrom [12]) that accommodation is the best response to entry and that entry is the best response to accommodation. This common knowledge, in turn, arises from the situation represented as a game of complete and perfect information in which all the firms are fully informed about the structure of the tree describing the game being played, about the payoffs accruing to all players, and about the others' past acts. As soon as the complete information assumption on the game is relaxed, so that the common knowledge condition no longer obtains, the logic of the backward induction breaks down. (This point is illustrated in Appendix B.) The possibility of actions taken in the past being a useful guide to future behavior in similar situations now opens up, and with this, reputations³ can come into

³ Most generally, a player's reputation in this context would be the beliefs that other players hold about his unknown characteristics and on the basis of which they predict his behavior. These beliefs would depend on their initial beliefs and on their observations of the player's past behavior. In the model we develop here, a simple sufficient statistic for these beliefs will be identified as the firm's reputation.

play. Further, once the lack of complete information gives rise to reputation possibilities, the players' equilibrium behavior will adjust markedly. The resultant equilibrium is then radically different than with complete information, for it has exactly the qualitative properties of Selten's intuitive solution. Practicing predation now gives one a reputation as a predator which is valuable in deterring entry. Thus, if entry occurs at an early stage, it will meet a predatory response, because any other response encourages further entry. Recognizing this, potential entrants at these early stages will enter only if the market is so lucrative that they are willing to face certain predation. Only as the horizon draws near and the number of markets which may still be entered declines will the firm be willing to share a market.

In the following sections we present our analysis of a version of Selten's model of multiple markets with sequential entry possibilities. In contrast to Selten, however, we allow that there is some doubt in the minds of the potential entrants concerning the established firm's options, motivations and behavior. We compute an equilibrium in this context which involves predation even by firms which find such a strategy to be costly in the short run, and we show that this equilibrium is the unique one involving sequential rationality by all players. We also investigate some of the major comparative statics properties of this model. The chief of these relate the value of a reputation—and the costs one will be willing to incur to obtain it—positively to the frequency with which the reputation may be used, as measured by the length of the horizon and the inter-period discount factor. The final sections presents a summary and some suggestions regarding both the implications of this analysis for policy and the possibilities for developing other formal models involving reputations.

2. THE MODEL

As suggested in the previous section, incomplete, asymmetrically distributed information plays a central role in our analysis of predation and, in particular, in undoing the logic leading to Selten's Chain Store Paradox. One obvious way to introduce the requisite informational asymmetry is to allow that the entrants do not know exactly the payoffs accruing to the established firm from an aggressive response to entry and that there is some positive probability that such a response is directly more profitable in any given stage than a response of peaceful coexistence. Kreps and Wilson [7] employ this approach, and we had also explored this avenue in earlier versions of this paper.⁴

⁴ While their analysis and ours originated independently, we have since benefitted greatly from having access to their ideas. They have also developed significant extensions of the "one-sided" reputation models studied here. Their paper is highly recommended to the reader.

Here we adopt a different approach in which we assume that the established firm definitely finds predation to be directly less profitable than sharing a given market. However, we also allow that there is some arbitrarily small, but nonvanishing, element of doubt in the minds of the entrants about whether their model of the established firm's options, motivation, and behavior is correct. In particular, our modeling allows that the entrants entertain some possibility that one or another simple behavioral rule guides the actions of the established firm. The nature of these rules is such that past behavior is repeated when similar circumstances arise. While yielding an equilibrium with exactly the same qualitative properties as that identified by Kreps and Wilson and by us in our earlier modeling, the present approach permits somewhat simpler arguments than we had needed before, it yields a strong uniqueness result, and it also appears to have some measure of applicability to a broad range of other problems. We will return to these issues of motivation, interpretation and justification after presenting the formal model.

We consider a game with $N + 1$ players. Player 0 is the established firm, while player n , $n = N, \dots, 1$, is a potential entrant in market n , where N is the first market threatened and 1 is the last. Associated with each player i is a random variable τ_i which is uniformly distributed on $[0, 1]$, independent of the other τ_j . We refer to a realization t_i of τ_i as the "type" of player i , $i = 0, 1, \dots, N$. As well, we have two strictly increasing, continuous functions, α , and β , where


$$\alpha: [0, 1] \rightarrow (-\infty, 0)$$

and


$$\beta: [0, 1] \rightarrow (-\infty, 1).$$

The function α will give the payoffs to firm 0 from preying at a particular stage as a function of its type, while β gives the payoff to an entrant which decides to stay out, again as a function of its type.


The $N + 1$ firms will play one of three possible games, each of which involves N repetitions of a particular stage game. (Only the established firm will know which of these actually obtains.) The first possibility is that the game is one where the n th stage is



The second possibility is that the game has as its n th stage



The third possibility is that the n th stage is



The first pattern is initially assessed a probability of $1/(1 + \epsilon + \delta)$, the second is given probability $\epsilon/(1 + \epsilon + \delta) > 0$, and the third is given probability $\delta/(1 + \epsilon + \delta) \geq 0$. We think of ϵ and δ as being small. Note that the realized value of τ_0 matters only when it is the first of these games that is being played. It is then convenient when referring to the situation where the second game is being played to abuse the notation by saying that $\tau_0 = \omega$ ($\omega > 1$) in this circumstance. Correspondingly, we will say that $\tau_0 = -\omega$ when it is the third game which is being played.

In each case, we have normalized the stage game payoffs so that the profit accruing to the established firm in any market in which it does not experience entry is 1 while its profit when peacefully sharing the market is 0, and so that the profit to the entrant is 0 if it meets predation and 1 if its entry elicits a nonaggressive response. These normalizations in no way effect the results: they solely serve to ease computations. In particular, they do not mean that the entrant "breaks even" when preyed upon, nor do they imply that, for example, the profits of the entrant if it is not preyed upon are equal to the profits of the established firm when there is no entry.

The specifications of the ranges of the α and β functions do have meaning in this context, however. That α is bounded above by 0, which is the profit payoff from preying, means that, other things being equal, any profit maximizing established firm would prefer to share a single market rather

than to prey. (We are thus building McGee's arguments into our formal model.) Similarly, the upper bound of 1 on β means that the entrants' outside opportunities are never better than sharing the market if entry will meet a passive response. The condition on α is used in our uniqueness argument, but may not be necessary; the condition on β serves only to simplify the arguments and could definitely be relaxed. Allowing that the realized values of β may be negative recognizes the possibility that entry might occur even if predation were certain. This could, for example, capture the idea that predation would fail against a particular entrant.

Our model is distinguished from Selten's in two major respects, both of which involve aspects of incomplete information. First, we assume that each player has a continuum of possible types, with different payoffs for each type. This assumption serves primarily to generate a pure strategy equilibrium. Taking α and β to be constant functions would eliminate this difference between our model and Selten's but would still generate a predatory equilibrium strategy for the established firm. Hence, the first distinction cannot account for the qualitatively different equilibrium behavior that we will find. The second difference is the positive ϵ and δ assumed in our model; in Selten's model $\epsilon = \delta = 0$. We have introduced positive values of ϵ and δ to capture the ideal that entrants entertain the possibility that a predatory response in one period might be part of a general aggressive pattern, and a cooperative response might be part of a general cooperative pattern. It is important to recognize that, if $\epsilon = \delta = 0$, there is no compelling reason for an entrant to suspect that any observed behavioral pattern might continue: past behavior, in that case, is utterly irrelevant in forecasting future behavior. It is precisely that irrelevance that leads to and is the heart of Selten's paradox. We shall see later that even as ϵ and δ approach zero, the potential entrants' strategies do not approach those specified by Selten: As ϵ , δ approach zero, the probability that predation deters entry can be bounded away from zero. However, the probability that predation actually occurs does converge to zero.

We assume that

$$\alpha(1) > \frac{-\rho(1 - \beta^{-1}(0))}{1 - \rho(1 - \beta^{-1}(0))} \equiv \alpha,$$

where $\rho \in (0, 1)$ is the discount factor used by firm 0. This condition turns out to be necessary for any reputation building to occur. While the basis for this claim must remain somewhat opaque for now, it will be shown that the right-hand side of the inequality represents the critical value of $\alpha(t_0)$ such that a firm of type t_0 is just indifferent about preying when the horizon is infinite and a single act of predation convinces the entrants that predation is certain in every future period. We also assume that $\beta(1) > 0$. Otherwise,

there is no possibility of deterring entry, since being preyed upon is always better than the alternative opportunity involved in not entering.

It is worth noting in this context that the assumption that the τ_i are uniform random variables on $[0, 1]$ in fact imposes no restriction on the model. This is because α^{-1} and β^{-1} are the cumulative distributions of the random payoffs.⁵ Further, although our assumptions that α and β are increasing functions means that there are no mass points in these distributions, this assumption on α is for convenience only. The strict monotonicity of β does assure some useful continuity in the optimal responses, but this too can be relaxed at the cost of complicating the equilibrium (see [7, Remark A]).

Initially, the structure of the game is common knowledge. This includes the values of δ and ϵ , the α and β functions, and the distributions of the random variables τ_i which determine the types. As well, it is common knowledge that only player i knows the value of τ_i and that only player 0 knows which of the three repeated stage games is being played. At each point in the game, each firm knows the history of the moves taken to that point by it and the other firms, but firm 0's payoffs in previous rounds are not observable by the other firms. Finally, it is also common knowledge that firm 0's payoff from the whole game is the present value of its profits at each stage, calculated with the discount factor ρ .

This framework corresponds to Harsanyi's treatment of games of incomplete information played by Bayesian players [5]. The distributions α^{-1} and β^{-1} reflect the various firms' beliefs about each other's payoffs, and the ϵ and δ reflect the doubts that the entrants have about whether their modeling of the established firm via the first repeated game is correct. Since the players are not sure about either the form of the tree describing the game they are playing (i.e., whether τ_0 is $+\omega$, $-\omega$ or in $[0, 1]$) or the payoffs accruing to various strategies (since the realized values of α and β are not public information), the situation is a game of incomplete information. Harsanyi's method for solving such games involves introducing a new game with complete but imperfect information, i.e., one in which the players all know the full game tree and all the payoffs but are not fully informed about the previous moves of the other players. In this game there is an additional player, Nature, which moves first. Nature plays a mixed strategy, selecting the actual types of the various players according to the probability distribution over types that describes the players' prior beliefs about one another. This move by Nature in our framework thus determines which of the three possible games is being played and the actual, realized values of the payoffs to the $N + 1$ firms. However, only firm 0 is informed about the

⁵ Throughout, for increasing functions $f: [0, 1] \rightarrow R$, we employ the conventions that if $x < f(0)$, then $f^{-1}(x) = 0$ and if $x > f(1)$, then $f^{-1}(x) = 1$.

outcome of Nature's choice of the game tree and about the value of τ_0 , and only firm n is informed about the realized value of τ_n . Thus, the game, while one of complete information, is also one with imperfect information, since a player's information sets typically include several different decision nodes, with the node that actually obtains having been determined by an aspect of Nature's move about which the firm is not informed. Harsanyi's approach is then to identify an equilibrium of this game with complete, imperfect information as the equilibrium of the original game of incomplete information.

In any game, a pure strategy for a player identifies an action to take in each information set. Here, a player's information sets are differentiated in part by its type. For example, given any stage n and any history of play by the firms up to that point, firm 0 still has a continuum of different information sets, one for each possible determination of its type through Nature's move. Thus, if only for this formal reason, a strategy for any firm must be a function of its type and thus must specify what it would do if its type were, say, some value t , even when its true (realized) type is something different.

This fact often seems to cause difficulties for those who have not previously dealt formally with games of incomplete information. One way to interpret the idea that a strategy must specify behavior for types that don't actually exist is to regard a strategy for a particular player as a conjecture in the minds of the others about its behavior. Then the dependence of the strategy on the player's type simply reflects the other players' making allowances in forming these conjectures that the player's behavior depends on its type, which is unknown to them. In our model, these conjectures by the entrants about firm 1's behavior are first about whether or not they have correctly modeled the established firm as choosing between predation and sharing, and secondly, if it does make such choices rather than reacting mechanistically, about the conditions under which it will prey. With this view, an equilibrium involves each player reacting optimally to its conjectures about the other's behavior, and these conjectures being consistent with the choices actually made.

Our solution of this game employs the concept of sequential equilibrium introduced by Kreps and Wilson [8]. This equilibrium notion requires that, at any decision node, the player take an action that maximizes its expected payoff, given its current beliefs and given that the others will henceforth follow the prescribed equilibrium strategies. These current beliefs (about, e.g., which node in an information set actually obtains) must be consistent with the player's initial beliefs, with any information that it may have available (directly or by inference) and, whenever possible, with the hypothesis that play has evolved to this point under the equilibrium strategies.

To define a sequential equilibrium formally, let H_n denote the history of the moves taken by the various firms from stage N to, but not including, stage n . Let H_n^+ denote the resulting history in stage $n-1$ when predation

occurred in stage n , let H_n indicate that sharing occurred, and let H_n^0 indicate that there was no entry. Let $\mathcal{H}_n$ denote the possible histories of play up to stage n . Then a strategy for firm 0 consists of N maps

$$s_0^n: \mathcal{H}_n \times \{-\omega\} \cup [0, 1] \cup \{\omega\} \rightarrow \{\text{Prey, Share}\},$$

$n = 1, \dots, N$, where $s_0^n(\cdot, \omega) = \text{Prey}$ and $s_0^n(\cdot, -\omega) = \text{Share}$. A strategy for entrant n is a map $r^n: \mathcal{H}_n \times [0, 1] \rightarrow \{\text{Enter, Stay Out}\}$. Then a sequential equilibrium is a strategy for each firm such that:

- (1) for each $n = 1, \dots, N$, each $H \in \mathcal{H}_n$ and each $t_n \in [0, 1]$,

$$r^n(H, t_n) = \begin{cases} \text{Enter} & \text{if } [1 - p_n(H)] > \beta(t_n), \\ \text{Stay Out} & \text{otherwise;} \end{cases}$$

- (2) for each $n = 1, \dots, N$, each $H \in \mathcal{H}_n$ and each $t_0 \in [0, 1]$,

$$s_0^n(H, t_0) = \begin{cases} \text{Prey} & \text{if } a(t_0) + \rho \bar{V}_{n-1}(t_0, H^+) > \rho \bar{V}_{n-1}(t_0, H), \\ \text{Share} & \text{otherwise,} \end{cases}$$

where the value function $\bar{V}$ is defined recursively, given $\bar{V}_0 \equiv 0$, by

$$\bar{V}_n(t_0, H) = q_n(H) \max[\rho \bar{V}_{n-1}(t_0, H^+), a(t_0) + \rho \bar{V}_{n-1}(t_0, H^+)] \\ + (1 - q_n(H)) [1 + \rho \bar{V}_{n-1}(t_0, H^0)];$$

and

- (3) for all n and all $H \in \mathcal{H}_n$

$$p_n(H) = \text{Prob}\{s_0^n(H, \tau_0) = \text{prey} | H\}$$

and

$$q_n(H) = \text{Prob}\{r^n(H, \tau_n) = \text{enter} | H\}.$$

In this definition, the p_n and q_n , which are interpreted as perceived probabilities of predation and of entry, respectively, represent the players' conjectures, and condition (3) is the rational expectations consistency requirement.⁶

Sequential equilibria are always Nash. Moreover, Kreps and Wilson have shown [8] that "for 'almost all' games [with finite strategy spaces], the

⁶ Note that the conditional expectations in (3) may not be well defined if H is a history of play that occurs with probability zero under the specified strategies. In this case there is some latitude possible in specifying $p_n(H)$ and $q_n(H)$, although not any arbitrary specification will do. See [7 and 13] for further discussion of this point in related contexts.

perfect and sequential equilibria 'nearly' coincide." Although the continuum of choices open here to Nature renders this result formally inapplicable to our model, the sequential equilibrium still clearly captures the basic idea behind perfectness. An advantage of the Kreps-Wilson equilibrium formulation over the perfectness approach is that it greatly eases the problem of computing and verifying that particular strategies constitute an equilibrium, since it allows us to use the methods of dynamic programming to analyze the players' decisions. To apply the methods of dynamic programming, it is necessary to define one or more state variables that summarize some of what the players know about the current position of the game. In the original Kreps-Wilson formulation, the beliefs of each player function as his personal state variable. In our model, the firm's reputation, which will be defined by a statistic that summarizes the history of play, will serve as a state variable.

3. EXISTENCE, UNIQUENESS AND PROPERTIES OF THE EQUILIBRIUM

To begin analysis of this game, note that at stage 1 (the last market), firm 0 must never prey as long as $t_0 \neq \omega$, since to do so lowers its payoff with no possible compensation. Thus, if it is known that $t_0 \neq \omega$, entry must occur at stage 1. But if firm 0 ever fails to prey, the entrants can all immediately infer that $t_0 \neq \omega$, (since $t_0 = \omega$ automatically yields predation). Selten's argument then applies, since it is common knowledge that firm 0 will not prey at stage 1. Thus, once it has ever failed to prey, firm 0 cannot gain by preying at stage 2, or, for that matter, at any other stage, since to do so simply squanders profits (α is negative) and cannot influence future entry decisions. Thus, if firm 0 ever fails to prey, it is clear that it will never prey again in equilibrium. It is also then clear that entry will occur in every succeeding market, since $\beta(t_n)$ is less than the payoff from unopposed entry for all t_n . Thus, in any sequential equilibrium, failure to prey against a particular entrant implies that the present value of the established firm's future payoffs from this and all succeeding stages is zero.

Now suppose that firm 0 has never failed to prey, but that entry has just occurred in market n . As just seen, failure to prey yields a value of zero to the established firm for the rest of the game. If, however, it adopts an alternative strategy which involves preying at stage n , then, given the strategies of the remaining entrants, it faces some list of possible patterns of future entries, predatory episodes, and unentered markets. In expected present value terms, let E , P , and M represent, respectively, the count of future entries, current and future predatory acts, and unshared markets resulting from the alternative strategy. Then the expected present value of future payoffs from following this strategy of predation is equal to $Pa(t_0) + M$, given our

normalizations. Since $P \geq 1$, this expression is increasing in t_0 , the value of τ_0 . Thus, if it pays firm 0 to prey at stage n when $\tau_0 = t$ (i.e., $Pa(t) + M > 0$), it also pays when $\tau_0 = t' > t$. In fact, then, the maximized value of future payoffs over all possible strategies is also increasing in t_0 , since it is the value of the maximum of increasing functions (including the constant function 0).

These points can be effectively illustrated graphically. The axes in Figure 1 are M and P , as defined above. Given the succeeding entrants' strategies, a particular choice of strategy at stage n by firm 0 results in a particular point in (P, M) -space. The convex hull of the resulting points is graphed. Note that $P + M \leq n$, and that the origin is always available. Firm 0's preferences over this space are given by the linear indifference curves corresponding to $Pa(t_0) + M = C$. Payoffs are increasing as we increase M or decrease P , and the slope of any particular type's indifference curves is $-a(t_0) > 0$. Thus, if $t_0 > t'_0$, then the t_0 indifference curves are flatter than those for t'_0 . It then is seen that the optimal P is at least weakly increasing in t_0 , and that if a given strategy yields a positive payoff for a particular value of t_0 , it yields a positive payoff for all higher values of t_0 .

Thus, in searching for an equilibrium, we can limit our attention to pure strategies for firm 0 which call for predation at a particular stage with a given history of play if and only if the value of τ_0 exceeds some critical value, where this value is 1 if the history to that point involves any failure to prey. (Clearly, for $\tau_0 = |\omega|$, there is nothing to specify. The reason we can restrict ourselves to pure strategies is that, with $Pa(t_0) + M$ increasing in t_0 , at most one single type would ever be indifferent between the two possible actions at any stage.)




FIG. 1. Monotonicity of Payoffs in τ_0 .

Thus, in equilibrium, if firm 0 has preyed whenever entry has occurred, the entrants can infer that either $\tau_0 = \omega$ or the value of τ_0 exceeds the maximum of the critical values governing its past predatory decisions. Let x be this maximum, and call x the *reputation* of firm 0 with the given history of play. We make the special conventions that if the established firm has ever failed to prey, its reputation is $x = -\infty$, and that if no entry has previously occurred, $x = -1$.

We will now characterize an equilibrium for this game in which at stage n :

- (1) firm n 's decision depends only on the value of τ_n and on firm 0's current reputation x ;
- (2) firm 0's decision (for the relevant case where $|\tau_0| \neq \omega$) depends similarly only on the value of τ_0 and on x ;
- (3) firm 0's new reputation on entering stage $n+1$ depends only his reputation entering stage n and the actions taken there.

This equilibrium involves N numbers, $x_1 \leq x_2 \leq \dots \leq x_N$, with the property that, so long as it has never previously failed to prey, firm 0 with $|\tau_0| \neq \omega$ will prey in response to entry in market n if and only if $\tau_0 > x_n$. Further, we show that $x_1 = 1$, that $x_n < x_{n+1}$ unless $x_{n+1} = 0$ or 1, that $\lim x_n = \max\{0, \alpha^{-1}(\alpha)\}$, and that, if $\alpha(0) > \alpha$, there exists a finite k such that, for all $n > k$, $x_n = 0$. Thus, in this equilibrium, the set of firms which will prey at stage n includes all those which will prey at any later stage when there are fewer markets to protect, and, for large N , any firm for which $\alpha(\tau_0) > \alpha$ will prey in the early rounds, regardless of the immediate cost. We also show how to compute the x_n values, demonstrate that this is the unique sequential equilibrium, and obtain some comparative statics results.

Since the entrants are to be looking only at their own types and firm 0's reputation, and since the revision of the reputation is to depend only on its current value and current actions, the expected present value to firm 0, when entering stage n , of playing optimally in this and all later stage depends only on the value t of τ_0 and on its current reputation x . Let $V_n(t, x)$ denote this value, and recall that, as argued above, V_n is increasing in t as the supremum of increasing functions. We will describe recursively a set of strategies based on these V_n functions, then verify that they do in fact constitute the unique sequential equilibrium.

Step 0. Initialization. Let $V_0(t, x) \equiv 0$ for all t and x . This is just an initialization based on a "dummy" stage 0. We also initialize by setting firm 0's reputation entering stage N at -1 .

Step 1. Reputation Revision at Stage n . Suppose firm 0's reputation entering stage n is $x \neq -\infty$. Then its reputation on entering stage $n+1$ is

Outcome in Market n	Reputation Entering Market $n+1$
No entry	x
Sharing	$-\infty$
Predation	$x \vee x_n$

where " $\vee$ " is the max operator and

$$x_n \equiv \inf\{x \in [0, 1] \mid t > x \text{ implies } \alpha(t) + \rho V_{n-1}(t, t) > 0\}.$$

If $x = -\infty$ at stage n , then firm 0's reputation at stage $n+1$ is also $-\infty$, independent of the actions taken at stage n .

Step 2. Firm 0's Actions at Stage n . With $\tau_0 = t \in [0, 1]$ and current reputation $x \neq -\infty$, we specify the firm 0 will prey at stage n if and only if

$$\alpha(t) + \rho V_{n-1}(t, x \vee x_n) > 0,$$

i.e., if and only if the current return plus the value of continuing optimally exceeds zero, which, as argued above, is the value of the payoff's resulting from failure to prey (i.e., $V_n(t, -\infty) = 0$). If $x = -\infty$, we specify that the firm will share the market if entry occurs.

This is clearly an optimal strategy for firm 0.

Remark: Suppose now that firm 0 does follow this strategy. What can subsequent entrants correctly infer from seeing firm 0 prey at stage n ? If the firm entered stage n with reputation x , then the subsequent entrants need consider only firm types t that exceed x . There are two possible cases.

First, if $x > x_n$, then $x = x \vee x_n$. In this case,

$$\alpha(t) + \rho V_{n-1}(t, x) \geq \alpha(x) + \rho V_{n-1}(x, x) > 0$$

so long as $t \geq x$, since α and V_{n-1} are increasing in t . The second inequality here follows from the definition of x_n and from $x > x_n$. Thus, if the firm has already demonstrated that it is of type $t > x_n$, it will surely prey at stage n .

Second, if $x \leq x_n < 1$, then $x_n = x \vee x_n$. Then the fact that α and V_{n-1} are increasing in t implies that

$$\alpha(x_n) + \rho V_{n-1}(x_n, x_n) \geq \alpha(t) + \rho V_{n-1}(t, x_n)$$

as $x_n \geq t$. Now, if $x_n > 0$, the continuity of V_{n-1} (to be verified shortly) implies that the left-hand side of the inequality is zero. Then the firm will in fact optimally prey if $t > x_n$, will optimally share if $t < x_n$, and is otherwise indifferent. If $x_n = 0$, then the left-hand side is non-negative, t is automatically greater than x_n , and predation is an optimal response, while if $x_n = 1$, then a similar argument shows the optimality of sharing. Thus, we

have verified that future entrants could correctly forecast that firm 0 will prey in market n with reputation $x \neq -\infty$ if and only if $r_0 > x \vee x_n$. This justifies the evolution of reputations that we specified.

Step 3. Firm n 's Actions at Stage n . In view of these computations, if $x > x_n$, then entrant n will in equilibrium correctly regard a predatory response as certain. If $x = -\infty$, entry will be correctly anticipated to meet a nonaggressive response. In general, we have

$p_n(x)$ = the probability in equilibrium that firm 0 will prey at stage n with reputation x , given that $r_0 \geq x$

= the probability assessed by entrant n that, given the reputation x , firm 0 will prey at stage n

$$p_n(x) = \begin{cases} \frac{\varepsilon + |1 - (x \vee x_n)|}{\varepsilon + (1 - x)} & \text{if } x \geq 0 \\ 0 & \text{if } x = -\infty \\ \frac{|\varepsilon + (1 - x_n)|}{|1 + \varepsilon + \delta|} & \text{if } x = -1. \end{cases}$$

The entrant's best response at n is then to enter if and only if

$$p_n(x) \cdot 0 + (1 - p_n(x)) \cdot 1 \geq \beta(t_n),$$

i.e., if and only if

$$t_n \leq \beta^{-1}(1 - p_n(x)) \equiv q_n(x).$$

Then $q_n(x)$ is precisely the probability that n enters at stage n when firm 0's reputation is x .

These strategies are clearly optimal responses to firm 0's strategy.

Step 4. Transition from Stage n to Stage $n + 1$. With the initialization for V_0 , the preceding steps allow us to calculate V_1 as $V_1(t, x) = 1 - q_1(x)$. To complete the recursive cycle, we now have to specify how to move from stage n to $n + 1$. For $x \neq -\infty$,

$$V_{n+1}(t, x) = q_{n+1}(x) \max\{0, a(t) + \rho V_n(t, x \vee x_n)\} + (1 - q_{n+1}(x))(1 + \rho V_n(t, x)).$$

The first term on the right-hand side is the probability of entry times the

maximum of the expected present value of the payoffs either from sharing the market and getting a reputation of $-\infty$ (where we have used the previously established fact that $V_n(t, -\infty) = 0$) or from preying now and then continuing optimally with the new reputation $x \vee x_n$. The second is the probability of no entry times the value of having this market as a monopoly and then continuing optimally with one's reputation unaffected.

Note that it is now clear that V_n is continuous if V_{n-1} is continuous. Thus, since $V_0 \equiv 0$ is continuous, all the V_n are.

It is also clear that V_n is the optimal value function and, thus, that the strategies do constitute a sequential equilibrium. Each firm is maximizing, given its beliefs and its conjectures about the other firms' strategies, the beliefs are consistent, and the conjectures are correct.

Before we turn to characterizing this equilibrium, some further discussion of the determination of the x_n values may be useful. In Fig. 2 we show $[0, 1] \times [0, 1]$ with the diagonal. Think of the horizontal axis as being reputations and the vertical as the type of the established firm. Then we ask: if preying at stage n were to yield a reputation x , which types of firms would be willing to prey? The set of such firms is $\{t \in [0, 1] \mid a(t) + \rho V_{n-1}(t, x) > 0\}$. Thus, the lowest type willing to prey is the inf over this set. Denote this inf as $h_n(x)$, and note that h_n is continuous and monotone decreasing. Then x_n is the unique fixed point of $h_n(x)$, i.e., the point where the graph of h_n crosses the diagonal. To see this, note that if $x \neq x_n$, then the set of firms willing to prey if it yields reputation x would not in fact justify the earning of that reputation.




FIG. 2. Determination of x_n .

We now seek to characterize the equilibrium. In so doing, the following function will be useful. Define $g: R^2 \rightarrow R$ by

$$g(t, v) = (1 - q_1(1))(t + pv) + q_1(1) \max(0, \alpha(t) + pv),$$

where $q_1(1)$ is the probability of entry against the "strongest" type of established firm in the last period or, equivalently, the probability of entry in any round against certain predation, i.e., $q_1(1) = \beta^{-1}(0)$. Observe that $V_{n+1}(t, 1) = g(t, V_n(t, 1))$ and for all t , g is increasing in v and

$$|g(t, v_1) - g(t, v_2)| \leq \rho |v_1 - v_2|,$$

so that $g(t, \cdot)$ is a contraction map. Let $\bar{v}(t)$ be the unique fixed point for each t . Then one can verify that

$$\begin{aligned} \bar{v}(t) &= \frac{1 - q_1(1)}{1 - \rho(1 - q_1(1))} & \text{if } t \leq x, \\ &= \frac{1 - q_1(1)(1 - \alpha(t))}{1 - \rho} & \text{if } t \geq x, \end{aligned}$$

where

$$x = \alpha^{-1} \left(\frac{-\rho(1 - q_1(1))}{1 - \rho(1 - q_1(1))} \right) = \alpha^{-1}(\alpha).$$

This fixed point will provide a bound on $V_n(t, \cdot)$.

Note, too, that the relationship $\alpha(1) > \alpha$ which was assumed earlier is simply equivalent to $x < 1$. (At this point, one can straightforwardly verify the claims and interpretations made in connection with this assumption.) Further, since g is an increasing contraction map, we have the following result.

LEMMA. If $v < \bar{v}(t)$, then $v < g(t, v) < \bar{v}(t)$.

Proof. Let $v < \bar{v}(t)$. Then

$$g(t, v) < g(t, \bar{v}(t)) = \bar{v}(t).$$

Also,

$$\begin{aligned} \bar{v}(t) - g(t, v) &= g(t, \bar{v}(t)) - g(t, v) \\ &\leq \rho(\bar{v}(t) - v) \\ &< \bar{v}(t) - v, \end{aligned}$$

so $v < g(t, v)$.

Q.E.D.

Note that $V_0(t, 1) = 0 \leq \bar{v}(t)$ (with strict inequality at least for $x < t$) and that $V_{n+1}(t, 1) = g(t, V_n(t, 1))$. Thus, in particular, we have $V_n(t, 1) \leq \bar{v}(t)$ for all n .

PROPOSITION 1. For any n , $x_{n+1} \leq x_n$, with strict inequality if $x_n \neq 0$ or 1. Further $x^* \equiv \lim_{n \rightarrow \infty} x_n = \max[0, x]$.

The proof of this result is given in Appendix C.

Proposition 1 establishes the key point that any firm (defined in terms of the costs it incurs in preying) which is willing to prey when n markets are still threatened by actual or potential entry will also prey when there are a greater number threatened. It further gives the asymptotic properties of the x_n . If $x < 0$, so that $x^* = 0$, then the convergence to x^* is necessarily in finite time. Thus, for all n sufficiently large, the established firm will prey for sure unless $\tau_0 = -\omega$. If $x > 0$, then the convergence is only asymptotic: $x_n > x$ for all n . Of course, for large enough values of n , the differences between x_n and x , between $V_n(t, x)$ and $\bar{v}(t)$, and between $p_n(x)$ and unity are arbitrarily small.

Note that none of these asymptotic properties depend on the particular values of $\epsilon > 0$ and $\delta \geq 0$.

In terms of Figure 2, the behavior of the h_n functions as n varies gives the evolution of the x_n values. If $\alpha(1)$ is sufficiently negative, no firm would find predation attractive if n is small, and the corresponding h_n functions are identically equal to 1. Eventually, however, acquiring a reputation of 1 (or close to it) by a single act of predation will become attractive for firms with high enough $\alpha(\tau_0)$ values, since such a reputation reduces the threat of entry to (essentially) $\beta^{-1}(0)$. Thus, at least near $x = 1$, $h_n(x)$ lies below x , and x_n becomes strictly less than 1. The correspondence between higher values of n and lower values of x_n then is a matter of h_n (at least near the diagonal) being above h_{n+1} . Finally, all the h_n converge in n to a constant function at $x^* = [0 \vee x]$, indicating that any type above x^* would be willing to prey to get such a reputation.

The equilibrium evolution of the game thus becomes clear. For large n , if $x < 0$, then once firm 0 has preyed and thereby revealed that it is not irrevocably committed to sharing ($\tau_0 \neq -\omega$), entrants correctly regard predation as certain. Thus, the only entrants which will attempt entry for large n after one predatory episode are those whose outside alternatives are so poor, relative to the profits available by entry, that they are willing to face certain predation. The probability of entry is then $\beta^{-1}(0)$, which may be zero. If $x > 0$, then a second act of predation is not absolutely certain (since $x_{n+1} > x_n > x$ for all n), but it is close to being so when n is large. Consequently, the probability of entry is close to $\beta^{-1}(0)$. In either case, the established firm sees predation at an early stage as leading to a long string of

monopoly returns, possibly interrupted by occasional predatory episodes. If its payoff while preying exceeds $a = a(x) = [-p(1 - \beta^{-1}(0))]/[1 - p(1 - \beta^{-1}(0))]$, it prefers this to the continuation value of zero which results from failure to prey. If $t < x$, sharing is preferred.

Before the first attempt to entry, if $\delta > 0$ it is still possible that the established firm is of the type that never preys. In this case, the probability of predation is strictly less than one, even when $x < 0$, and so one would expect (with probability strictly greater than $\beta^{-1}(0)$) to see a test of firm 0's fortitude by an entrant with a relatively low value of τ_n and thus low opportunity costs in entering. If this entry meets predation, then the firm's reputation immediately jumps to $x \geq x^*$, and the game proceeds as described above.

As the number of markets remaining threatened decreases over time and the horizon approaches, x_n increases. This begins immediately if $x > 0$, and at some finite date if $x < 0$. Consequently, the conditional probability that $\tau_0 > x_n$, given the current reputation x , and thus the probability of predation will start to decrease. If at some stage this probability of predation is sufficiently low relative to $\beta(t_n)$, entry will occur. If $t \leq x_n$, entry will meet a nonaggressive response. The possibility that $\tau_0 = \omega$ is then eliminated, it thus becomes certain that predation will not occur in market 1. Selten's logic takes over, and unopposed entry occurs in all remaining markets. If $t > x_n$, predation occurs. The established firm's reputation now rises to x_n , the probability of predation at stage $n-1$ now is $p_{n-1}(x_n) > p_{n-1}(x)$, the attractiveness of further entry is reduced from what it would have been if the value of τ_0 were only known to exceed x rather than x_n , and the probability of entry in future rounds correspondingly falls.

We will now demonstrate that this is the unique sequential equilibrium.

We have already observed that, in any sequential equilibrium, if the established firm ever fails to prey, then there must be entry and sharing at every later stage. Thus, it suffices to consider only what happens when there has been no sharing. We have also observed that at any stage n with any history of play to that point, firm 0 with $|\tau_0| \neq \omega$ will prey if and only if τ_0 exceeds some critical value. Thus, at any sequential equilibrium all that can be inferred from a given history involving no sharing is that $\tau_0 = \omega$ or $\tau_0 \in (x, 1]$ for some x . As well, we have also observed that we need consider only pure strategies.

We now establish inductively that any sequential equilibrium must agree with the one we have described earlier at all stages $m \leq n$. The induction is made on n , and the case $n = 0$ is immediate. Fix any sequential equilibrium and suppose that the result holds for $n = k$. Then the value of entering stage k with $\tau_0 = t$ and reputation x is $V_k(t, x)$. Now consider firm 0 in stage $k+1$, where the history is H and the corresponding reputation is $\bar{x}$. Let $x \geq \bar{x}$ denote the reputation that 0 would have at this equilibrium if entry and

predation occurred at stage $k+1$. Then firm 0 of type $\tau_0 = t$ would choose to prey or not according as

$$a(t) + \rho V_k(t, x) \geq 0.$$

The reputation x must be consistent with 0's actual strategy at equilibrium, so we must have for $t \geq \bar{x}$ that

$$a(t) + \rho V_k(t, x) \geq 0 \quad \text{as } t \geq x.$$

Using the monotonicity of V_k , it is straightforward to check that the unique $x \geq \bar{x}$ for which this holds is $x = \bar{x} \vee x_k$. Hence, the conditional probability that entry will meet predation at stage $k+1$ given the history H and reputation $\bar{x}$ is $p_{k+1}(\bar{x})$. The unique best response to this for firm $k+1$ is to enter if and only if $\beta(t_{k+1}) < p_{k+1}(\bar{x})$, and this completes the induction. Notice that the fact that strategies depend on history only through reputations has been derived, not assumed. What the induction establishes is that if history affects play only through the reputation for stages $m \leq k$, then the same is true at stage $k+1$. Since this is trivially true for $k = 0$, it is true for all k .

The comparative statics properties of this equilibrium can be stated rather simply. Those on N have already been established: the values of the x_n are, in fact, independent of N , so that the probabilities of entry and of predation in round $n < N$ do not depend on N . Of course, we have already shown that x_n converges monotonically down to $\{0 \vee x\}$. Increases in ρ , the discount factor, tend to lower the x_n 's. To see the sense behind this claim, note that against given strategies by the entrants, a higher value of ρ increases (in present value terms) the reward to preying. The costs that the established firm would be willing to incur in preying then increase, lowering the cut-off level of x at which predation occurs. However, such a lower value for x_n will make entry less attractive, lowering $q_n(x)$ and increasing the benefit from predation at stage $n+1$. Thus we have a snowballing effect, with everything working in the same direction. (A formal proof of all the comparative statics results we will claim here is given in Proposition 2.)

Together these two results indicate that the value of a reputation—and the costs one would incur to achieve it—increase with the frequency with which it may profitably be used. Being able to use the reputation either more often at a given rate (increases in N) or with less delay between uses (increases in ρ) increases the incentives for building and maintaining a reputation. We will later suggest some interesting implications that would follow if this intuitively appealing result carries over to other models.

The effects of changing the α or β functions in certain ways can also be identified. Adding any positive function to β (while still maintaining

monotonicity and the upper bound of 1) yields a new function whose inverse, as the distribution of the outside opportunities, stochastically dominates the original distribution β^{-1} . This change lowers the probability of entry against any given reputation, given the established firm's strategy. However, the lower values of the q_n tend to increase the value of preying (since they mean that predation is more likely to deter future entry), so again we have the snowballing effect we saw in the case of ρ . Similarly, an increase in α , which shifts α^{-1} in the sense of stochastic dominance, also decreases the x_n . Thus both an increase in the outside profit opportunities of the entrants (which lowers the attractiveness of entry) and a decrease in the costs of preying tend to lower the x_n , thereby increasing the probability of predation with any given reputation and lowering the probability of entry.

Finally, a higher value of ϵ , the probability put on an automatically aggressive response, also works in the same way to lessen the x_n , and thus to increase the probability of predation with a given reputation and lower the probability of entry.

All these properties are established formally by the following result, which is proven in Appendix D.

PROPOSITION 2. Suppose $\rho \geq \bar{\rho}$, $\alpha \geq \bar{\alpha}$, $\beta \geq \bar{\beta}$ and $\epsilon \geq \bar{\epsilon}$. Then $x_n \leq \bar{x}_n$ for all n .

In this modeling, the possibility that $\tau_0 = -\omega$ (with probability δ) was included largely for the sake of symmetry: either form of mechanistic behavior pattern is possible. The sole effect of δ on the equilibrium is that it encourages an early test of firm 0 by some entrant. In particular, the x_n values do not depend on the value of δ . In fact, as long as the probability that $\tau_0 = \omega$ is kept bounded from zero, we could also have allowed that the entrants also entertained all the theories about the established firm's behavior which involve it preying against every entrant until there are exactly k markets left, $k = 1, 2, \dots$, and then never preying again. The equilibrium would not have been qualitatively changed.

Finally, let us examine the behavior of the equilibrium as $\epsilon \rightarrow 0$ and thus as we approach the original Selten model. Consider the case $N = 2$: the arguments can be extended for any N . By Proposition 2, x_1^* is increasing as $\epsilon \rightarrow 0$ and so has a limit x_1^0 . Then by the definition of x_1^* , we have $q_1^*(x_1^*) \rightarrow \max(0, 1 + \alpha(x_1^0)/\rho)$. Since α is bounded above by 0, we then have that the limit is strictly less than one. Further, if $x_1^0 \neq 1$, then we have $\lim q_1^*(x_1^*) = \lim \beta^{-1}[1 - \epsilon/(\epsilon + 1 - x_1^*)] = 1$, and this contradiction shows that $x_1^0 = 1$. Thus, as $\epsilon \rightarrow 0$, the set of firms which will prey in equilibrium becomes null. However, since $x_2^* = 1$, predation at stage 2 deters entry with probability equal to $\min[1, -\alpha(1)/\rho] > 0$. Note that this a failure of upper hemicontinuity of the sequential and perfect equilibrium correspondence as we

change the information structure: the established firm's equilibrium strategy converges to that identified by Selten, but the entrants' strategies do not.

Another way to express the striking importance of even very small positive values of ϵ is to consider how the deterrent effect of a strategy of always preying depends on ϵ . The easiest case to consider is when $\beta(0) > 0$. Suppose entry actually occurs at k dates $m_1, \dots, m_k$, where m_i occurs after m_{i-1} , and these entries are always met by predation. Then by examining the entrants' strategies we find that for any two successive entry dates m_i and m_{i+1} , $1 - \beta(0) > 1 - \beta(\tau_{m_i}) > [\epsilon + 1 - x_{m_i}]/[\epsilon + 1 - x_{m_{i+1}}]$, where $x_{m_{i+1}} \equiv 0$. This gives a set of k inequalities which when multiplied together lead to $[1 - \beta(0)]^k > [\epsilon + 1 - x_{m_1}]/[\epsilon + 1 - x_{m_k}] > \epsilon/(1 + \epsilon)$. Hence, $k < \ln[\epsilon/(1 + \epsilon)]/\ln[1 - \beta(0)]$. This provides an upper bound on the number of entries that can arise for any realizations of $\tau_1, \dots, \tau_k$ when entrants follow their equilibrium strategies and firm 0 adheres to an aggressive strategy. Note that the bound is independent of N and grows only logarithmically in ϵ . Even tiny values of ϵ can lead to moderate bounds on k . For example, if $\beta(0) = 0.75$, then for $\epsilon = 10^{-3}$ one obtains $k \leq 4$, for $\epsilon = 10^{-4}$ one obtains $k \leq 9$, and for $\epsilon = 10^{-5}$, $k \leq 14$. Thus, tiny elements of uncertainty can produce a significant deterrence of entry.

4. SUMMARY AND CONCLUSIONS

We have demonstrated that the presence of informational asymmetries can lead a firm operating in several markets to adopt a predatory strategy against entrants, even though such behavior is irrational when viewed in the context of a single market in isolation and even though there are only a finite number of potential entrants. This same point is made in the one-sided uncertainty model in the accompanying paper by Kreps and Wilson [7]. We view this model and our work as complementary, in that they display two different ways in which the recognition of informational asymmetries can "resolve" the Chain-Store Paradox. This resolution comes about because the informational asymmetry gives the entrants reason to forecast future actions on the basis of past behavior. This in turn gives the established firm reason to prey in order to build a reputation which leads future entrants to predict that they too are likely to meet predation.

The particular asymmetry we have introduced involves the entrants being less than certain that they are correct in modeling the established firm as rationally choosing between predation and peaceful coexistence. Specifically, we allow that they entertain the possibility that an episode of predation—which ought never to occur in the complete information set-up—may be part of a general pattern of predatory behavior. The recognition of this possibility

then results in the marked change in the equilibrium behavior that we have demonstrated.

There are numerous reasons why this element of uncertainty should exist. On the one hand, the entrants could be not completely sure about the game being played. For example, it might be that the established firm could actually be involved in some bigger game of which the one studied here is only a part and that the firm's equilibrium strategy in this larger game calls for it to prey in these markets. A second possibility is that in the game actually being played, the established firm may be able to precommit itself to an aggressive course of action and may have done so. Other scenarios involve the entrants allowing that the firm is not behaving as a fully rational game theorist.

For example, the fact that firms involve many individuals, each with his or her own preferences and information, suggests that the appropriate model of the firm would be one of group decision making, and there is no compelling reason for choices in such situations to correspond to the maximization of a single utility function. Alternatively, the firm may have well-defined utility function, but it may not calculate fully the equilibrium in the game being played. Instead, it may have some more or less arbitrary conjectures about how the entrants will behave in response to its actions, and its preying is the optimal behavior given these conjectures. In fact, Rosenthal [14] has argued that games of perfect information ought to be analyzed as if the players were decision theorists in this sense, rather than game theorists whose conjectures about others' behavior must be correct. In this context, Rosenthal suggested that his approach could resolve the Chain-Store Paradox, and Macgregor [9] has verified this conjecture. Thus, for example, if the established firm believed that the entrants would simply forecast that past behavior would be repeated, it would be led to prey.

In this regard, Scherer's discussion ([15, pp. 338-339] of predation as a deterrence strategy is of interest. He notes that if entrants perceive different markets as sufficiently similar and if "business managers extrapolate from past events rather than sizing up the probabilities in each new situation", then predation may have a deterrent effect. He goes on to note that very little is known about the extent to which such extrapolation occurs. Our analysis shows that even if everyone is fully rational and does "size up the probabilities in each new situation," as long as there is some doubt about whether this is the case, predation can emerge.

The implications of this analysis for antitrust policy are straightforward. In multiple market situations, predation can be a rational strategy which deters entry and thus supports monopoly. Thus, any tendency to discount the likelihood or significance of predation on the basis of its presumed irrationality should be checked when there are multiple markets which might reasonably be regarded by potential entrants as similar. (An interesting

paper by Easley *et al.* [4], in which firms can make their markets seem to be intrinsically unprofitable candidates for entry by predation, yields similar conclusions.) Firms serving several geographically distinct markets are only one case where the model and its conclusions might apply. Firms with broad product lines or those in which on-going technological change yields a pattern of new product introduction over time are others. In this context, it is worth noting explicitly that predation will only rarely need to be practiced. The credible threat of predation will deter all but the toughest entrants (those with low values of t_n), and so the occasions when the firm will be called upon to carry out its threat will be infrequent.

Two factors in our model lead to the emergence of reputations: the informational asymmetries and the repeated actions with the possibility of observing past behavior. These conditions will be necessary for reputation building to occur in general, and it would further seem that they are sufficient: in any situation where individuals are unsure about one another's options or motivation and where they deal with each other repeatedly in related circumstances (or where past dealings with other people are observable), we would expect to see reputations develop. A clear example of such a situation is in problems involving the choice of product quality. Other examples in economics arise in credit relationships, in labor negotiations and strikes, in implicit contract models, and in the provision of auditing services, bond ratings, job recommendations and the like. The reader can easily provide more examples. Moreover, applications outside the traditional bounds of economics are at least as numerous. We expect that the methods of this paper can be applied in each of these contexts to yield important insights on the nature of observed behavior. We also would expect that the main comparative statics result of this paper—that the value of a reputation and the extent of reputation building increase with the frequency of the opportunities for its use—will prove generally true. If so, this would provide insight on such issues as the costs of social mobility in terms of reducing the opportunities and incentives for building reputations for honest behavior, quality service, and the like.

The particular way in which we have introduced the informational asymmetry in this paper is to assume that individuals ascribe some positive probability to there being some alternative theory which determines behavior. In some situations, there will be an alternative theory which is particularly natural. In others, however, many perturbations of the "basic" complete information model will exist, and the choice between them may appear arbitrary. This suggests the need for a notion of robustness for these alternative theories. One such notion is that a theory would be robust if its presence led to behavior which did not refute the theory. We plan to explore such ideas in the context of developing reputation models of some of the phenomena discussed above.

APPENDIX A

In this Appendix we examine an infinite horizon version of Selten's chain-store model. As we demonstrate, one trivially establishes existence of an equilibrium where any attempted entry is met by predation, even with complete information. This equilibrium then resembles that in Section 2 for large N . However, in contrast with the finite horizon, incomplete information model, the set of equilibria in this model is large and contains many intuitively unappealing strategy combinations.

Let Π_0^P , Π_0^C and Π_0^M denote the incumbent's payoffs to predation, coexistence and monopolization in a single market, and let the corresponding payoffs to any entrant be Π_1^P , Π_1^C and Π_1^M . Given a discount factor ρ for the incumbent, and denoting the play at round n by h_n , where $h_n = +1$ denotes predation, $h_n = 0$ denotes no entry and $h_n = -1$ denotes sharing, the incumbent's payoff is

$$\Pi_0^C \sum_0^{\infty} \rho^n I_{h_n = -1} + \Pi_0^C \sum_0^{\infty} \rho^n I_{h_n = 0} + \Pi_0^M \sum_0^{\infty} \rho^n I_{h_n = +1}$$

where I_A is the indicator function for the event (set) A , i.e., $I_{h_n = a}$ is equal to one for those n such that $h_n = a$ and zero otherwise.

We assume that the game is one of perfect information, so that all players know the structure of the game tree, the payoffs, and the actions taken all previous moves. In contrast to the set-up in the text, number markets forward in time: market n opens before market $n+1$. Assume $\Pi_0^P < \Pi_0^C < \Pi_0^M$ and $\Pi_1^P < \Pi_1^C < \Pi_1^M$.

We now claim that, if $\Pi_0^P + \Pi_0^M \sum_1^{\infty} \rho^n > \Pi_0^C \sum_0^{\infty} \rho^n$, then the following strategies constitute a sequential equilibrium:

Firm 0: if $h_1, \dots, h_{n-1} \geq 0$ and entry occurs in market n , then prey.
If $h_k = -1$ for any $k < n$ and entry occurs in market n , then share this market.

Firm n : If $h_1, \dots, h_{n-1} \geq 0$, do not enter.
If $h_k = -1$ for any $k < n$, enter

So long as $\Pi_0^P + \Pi_0^M \sum_1^{\infty} \rho^n > \Pi_0^C \sum_0^{\infty} \rho^n$, not only do these strategies have the mutual best-response property that makes them Nash equilibria, but also this property obtains for the appropriately truncated strategies starting from any node in the game tree, so in fact they yield a sequential equilibrium.

However, as is again easily shown, the strategies of never preying for the established firm and all potential entrants always entering, regardless of the history, also constitute a perfect equilibrium. Moreover, if $\Pi_0^P + \Pi_0^M \sum_1^{\infty} \rho^n >$

$\Pi_0^C(1 + \sum_1^{\infty} \rho^n)$, then there are uncountably many pure-strategy perfect equilibria; indeed, there are uncountably many of the following form, parameterized by S .

Firm 0: If conditions (a), (b) and (c) hold at stage n , then prey at n :

- (a) $h_k \geq 0$ for all $k < n, k \in S$
- (b) $h_k \leq 0$ for all $k < n, k \notin S$
- (c) $n \in S$.

Otherwise, do not prey at n .

Firm n : If (a), (b), and (c) hold, stay out. Otherwise, enter.


Among the sets S that make these strategies a perfect equilibrium are: $S = \{k | k \geq i\}$ for any choice of i , and $S = \{k | k \text{ is divisible by } i\}$, for any sufficiently large integer i . In view of this multiplicity, perfect equilibrium lacks any effective predictive power in this infinite horizon model. Furthermore, predation would never be observed in the pure-strategy equilibria of this model, so it also fails to explain apparent observed behavior.

For these reasons, and because we also believe that the finite horizon model has its own inherent interest, we concentrate on it in the text.

APPENDIX B

The following example illustrates our claim regarding the role of lack of common knowledge in generating predation.

Specifically, we consider a situation with the established firm, 0, and two entrants, 1 and 2. There are three possible states of the world, a , b , and c , all of which are equally likely. In state a , the game at each stage is such that entry results automatically in predation. In states b and c , the stage game is



We can think of state a as one where the established firm is a predator.

The firms differ in their information, which can be described by the following partitions:

Firm 0: $\{\{a\}, \{b, c\}\}$

Firm 1: $\{\{a, b\}, \{c\}\}$

Firm 2: $\{\{a\}, \{b\}, \{c\}\}$

Note that firm 2 has perfect information, and that in state c it is common knowledge between the two entrants and also between the established firm and entrant 2 (who is the first to threaten entry) that firm 0 is not a predator. Yet in equilibrium, firm 2 is credibly threatened by predation! To see this, one need only note that the following constitute a set of sequential equilibrium strategies:

Firm 2: Stay out

Firm 1: Enter in event $\{c\}$, or if 2 entered and 0 failed to prey. Otherwise, stay out.

Firm 0: Prey if firm 2 enters. If firm 1 enters, prey in event $\{a\}$, and do not prey in $\{b, c\}$.

What makes this work is that, when the actual state is b or c , firm 0 does not know whether 1 knows that it (0) is not a predator, and everyone knows this. Note that 0 must take the same act in events b and c , since it cannot distinguish them. Thus, it can prey against firm 2 either in both states or in neither. If it shares in both states, it gets an expected return of 0. That this will result when c is the actual state is clear. When b obtains, failure to prey against 2 allows 1 to infer that it will not meet predation, and so entry in the last market again follows. If firm 0 preys in both states, then it absorbs a cost of -0.2 against firm 2. When c is the state, it still experiences entry, so its total payoff is -0.2 . In state b , however, the act of predation prevents 1 from inferring whether state a or b obtains. Entry in state a yields 1 a payoff of 0, entry in b yields it 1, and since it views the two as equally likely, it prefers to stay out and receive 0.8. Thus, from 0's point of view, predation yields equal chances of $-0.2 + 1$ and $-0.2 + 0$, for an expected return of 0.3. Since this exceeds the expected return of zero resulting from failure to prey, it preys. Given that 0 will prey no matter what the state, firm 2 is then deterred from entering. This is true even when (in state c) it and firm 0 both know that sharing is directly more profitable for 0, both know that both know this, both know that both know that both know this, *ad infinitum*.

In [6], this type of modeling is used to explain observed behavior in the finitely repeated prisoners' dilemma game.

APPENDIX C

In this Appendix we establish Proposition 1. The first part of the argument involves a rolling induction on the following list of propositions.

(P0) $V_n(t, x)$ is continuous in (t, x) and nondecreasing in t .

(P1) $V_n(t, x)$ is nondecreasing in x .

(P2) $x_{n+1} \leq x_n$.

(P3) For all u and all $x \leq x_{n+1}$,

$$1 + \rho V_n(t, x) \geq \max[0, u(u) + \rho V_n(t, x_{n+1})].$$

Since we have already established (P0), we turn to proving (P1)–(P3).

Proof of (P1)–(P3). We proceed by a "rolling induction" on n . The initialization for $n = 0$ is straightforward if we define $x_0 \equiv 1$.

Now suppose (P1) to (P3) hold for all $n \leq m$, and let us check the case for $n = m + 1$.

(P1): Rearranging the recursion for V_{m+1} yields:

$$\begin{aligned} V_{m+1}(t, x) = & \max(0, a(t) + \rho V_m(t, x \vee x_{m+1})) \\ & + (1 - q_{m+1}(x))[(1 + \rho V_m(t, x)) \\ & - \max(0, a(t) + \rho V_m(t, x \vee x_{m+1}))]. \end{aligned}$$

The max term, the $(1 - q_{m+1}(x))$ term, and the bracketed term are all non-negative (by the inductive hypothesis (P3)) and non-decreasing in x (by the inductive hypothesis (P1)). Thus, we have (P1) for $n = m + 1$.

(P2): By the just proven fact that (P1) holds for $n = m + 1$, we have for all x that $V_{m+1}(x, x) \leq V_{m+1}(x, 1)$. Together with the observation following the lemma, this gives $V_{m+1}(x, x) \leq \bar{v}(x)$. So, using the lemma and the fact that $V_{m+1}(x_{m+1}, x_{m+1}) = g(X_{m+1}, V_m(x_{m+1}, x_{m+1}))$, (which is seen to be just the defining recursion once we note that $q_1(1) = \beta^{-1}(0) = q_{m+1}(x_{m+1})$),

$$\begin{aligned} a(x_{m+1}) + \rho V_{m+1}(x_{m+1}, x_{m+1}) \\ = a(x_{m+1}) + \rho g(x_{m+1}, V_m(x_{m+1}, x_{m+1})) \\ \geq a(x_{m+1}) + \rho V_m(x_{m+1}, x_{m+1}). \end{aligned}$$

If $x_{m+1} = 1$, there is nothing to prove. If $x_{m+1} < 1$, then the right-hand side above is at least zero (by definition of x_{m+1} and continuity of V_m). Hence, by

the monotonicity and continuity of α and V_{m+1} and by the definition of x_{m+2} , we have $x_{m+2} \leq x_{m+1}$. (As we argue in the main text, this inequality will turn out to be strict if $x_{m+1} > [0 \vee x]$.)

(P3): We now must show that (P3) holds for $n = m + 1$. The case where $0 \geq \alpha(t) + \rho V_{m+1}(t, x_{m+2})$ is seen immediately to follow from the nonnegativity of $V_{m+1}(t, x)$. This is turn results from the availability of the strategy of always sharing, which yields a non-negative payoff. Thus we concentrate on the case $0 < \alpha(t) + \rho V_{m+1}(t, x_{m+2})$. Using the recursive expression for V_{m+1} , we then have

$$\begin{aligned} & 1 + \rho V_{m+1}(t, x) - \alpha(t) - \rho V_{m+1}(t, x_{m+2}) \\ &= 1 + \rho [q_{m+1}(x) \max\{0, \alpha(t) + \rho V_m(t, x_{m+1})\} \\ &\quad + [1 - q_{m+1}(x)] [1 + \rho V_m(t, x)] \\ &\quad - \alpha(t) - \rho [q_{m+1}(x_{m+2}) \max\{0, \alpha(t) + \rho V_m(t, x_{m+1})\} \\ &\quad + [1 - q_{m+1}(x_{m+2})] [1 + \rho V_m(t, x_{m+2})]] \\ &= \rho [q_{m+1}(x_{m+2}) - q_{m+1}(x)] [1 + \rho V_m(t, x) - \max\{0, \alpha(t) + \rho V_m(t, x_{m+1})\}] \\ &\quad + \rho [1 - q_{m+1}(x_{m+2})] [\rho V_m(t, x) - \rho V_m(t, x_{m+2})] + 1 - \alpha(t). \end{aligned}$$

By the induction hypothesis (P3), the first term in braces is nonnegative. Since $x_{m+2} \geq x$ by hypothesis and since q_{m+1} is decreasing in x (as is easily verified), we have $-1 \leq q_{m+1}(x_{m+2}) - q_{m+1}(x) \leq 0$. Since V_m is increasing in x by the induction hypothesis and since $x \leq x_{m+2}$, the second term in braces is non-positive. Thus, replacing $\rho [q_{m+1}(x_{m+2}) - q_{m+1}(x)]$ by -1 and $\rho [1 - q_{m+1}(x_{m+2})]$ by its upper bound, 1, we obtain, for $x \leq x_{m+2}$,

$$\begin{aligned} & 1 + \rho V_{m+1}(t, x) - \alpha(t) - \rho V_{m+1}(t, x_{m+2}) \\ &\geq -[1 + \rho V_m(t, x) - \max\{0, \alpha(t) + \rho V_m(t, x_{m+1})\}] \\ &\quad + [\rho V_m(t, x) - \rho V_m(t, x_{m+2})] + 1 - \alpha(t) \\ &= \max\{0, \alpha(t) + \rho V_m(t, x_{m+1})\} - [\alpha(t) + \rho V_m(t, x_{m+2})], \end{aligned}$$

which we must show to be non-negative. For this, it is sufficient that

$$V_m(t, x_{m+2}) \leq V_m(t, x_{m+1}).$$

But this inequality follows from (P1) and (P2), so we are done.

Thus, in particular, $x_{n+1} \leq x_n$. We now establish the remaining claims made in the statement of the Proposition.

Since $\{x_n\}$ is a bounded, nonincreasing sequence, it has a limit $x^* \geq 0$. If the firm ever preys (without having failed to do so before), it acquires a reputation $x \geq x^*$. For any such x and any n ,

$$\begin{aligned} p_n(x) &= \frac{\varepsilon + 1 - (x \vee x_n)}{\varepsilon + 1 - x} \\ &\geq p_n(x^*) = \frac{\varepsilon + 1 - (x^* \vee x_n)}{\varepsilon + 1 - x^*}. \end{aligned}$$

But since $x_n \rightarrow x^*$, this last term converges to 1 as $n \rightarrow \infty$. Thus, $p_n(\cdot)$ converges uniformly for $x \geq x^*$ to 1 as n goes to infinity, and so $q_n(x) = \beta^{-1}(1 - p_n(x))$ converges uniformly to $\beta^{-1}(0) = q_1(1)$. Predation is certain in the limit, and only those willing to be preyed upon will enter.

Then, since

$$\begin{aligned} V_n(t, x) &= q_n(x) \max\{0, \alpha(t) + \rho V_{n-1}(t, x \vee x_n)\} \\ &\quad + (1 - q_n(x)) [1 + \rho V_{n-1}(t, x)] \end{aligned}$$

for $\rho < 1$ we must have (by standard results in discounted dynamic programming)

$$\lim_{n \rightarrow \infty} V_n(t, x) = V_\infty(t, x) \quad \text{for } x \geq x^*,$$

where $V_\infty(t, x)$ is the unique solution to

$$\begin{aligned} V_\infty(t, x) &= q_1(1) \max\{0, \alpha(t) + \rho V_\infty(t, x)\} \\ &\quad + [1 - q_1(1)] [1 + \rho V_\infty(t, x)] \\ &\equiv g(t, V_\infty(t, x)). \end{aligned}$$

But, by definition, this solution is the fixed point, $\bar{v}(t)$. Thus, in the limit, predation gives the firm a reputation $x \geq x^*$ and a continuation value $\bar{v}(t)$. Then, the limiting condition guaranteeing that the firm will prey is $\alpha(t_0) + \rho \bar{v}(t_0) > 0$, i.e., $t_0 > x$. Thus, $x^* = \lim x_n = [0 \vee x]$.

Since we assumed $x < 1$, we have thus demonstrated that predation which is costly in the short run will still occur for n large enough. Reexamination of the argument establishing (P2) now in fact reveals that $x_n < x_{n-1}$ unless x_{n-1} is 0 or 1, since $\bar{v}(t)$ is strictly greater than $V_n(t, x)$ for all the relevant values of t . Q.E.D.

APPENDIX D

This proof of Proposition 2 involves a rolling induction on n based on the following two statements.

$$(P4) \quad \forall n, \quad x_n \leq \bar{x}_n$$

$$(P5) \quad \forall x \leq x_n, \quad V_n(x, x) \geq \bar{V}_n(x, x) \geq 0.$$

Proof. Notice that (P5) for $n = k$ implies (P4) for $n = k + 1$. To see this, note that if k were such that $x_{k+1} > \bar{x}_{k+1} \geq 0$, then, using the induction hypothesis (P5), the monotonicity of α and V_k , and the fact that $x_{k+1} > 0$, we have

$$\bar{\alpha}(\bar{x}_{k+1}) + \bar{\rho}\bar{V}_k(\bar{x}_{k+1}, \bar{x}_{k+1}) < \alpha(x_{k+1}) + \rho V_k(x_{k+1}, x_{k+1}) \leq 0,$$

which contradicts the definition of $\bar{x}_{k+1}$ and the continuity of $\bar{\alpha}$ and $\bar{V}_k$.

To show (P5), note first that both (P4) and (P5) hold trivially for $n = 0$ and that non-negativity was shown in the proof of Proposition 1. Suppose (P5) holds for $n = k$ and (P4) holds for $n = k + 1$. Then, for $x \leq x_{k+1}$,

$$V_{k+1}(x, x) = q_{k+1}(x) \cdot 0 + (1 - q_{k+1}(x))(1 + \rho V_k(x, x)),$$

$$\bar{V}_{k+1}(x, x) = \bar{q}_{k+1}(x) \cdot 0 + (1 - \bar{q}_{k+1}(x))(1 + \bar{\rho}\bar{V}_k(x, x)),$$

where the zero payoff in the first expression is because firm 0 will not prey if $x \leq x_{k+1}$, and the second follows since $x \leq x_{k+1} \leq \bar{x}_{k+1}$. For the claimed inequality, it now suffices, given (P5) for $n = k$, to show that $q_{k+1}(x) \leq \bar{q}_{k+1}(x)$. But $x_{k+1} \leq \bar{x}_{k+1}$ and $\bar{\epsilon} \leq \epsilon$ implies $\bar{p}_{k+1}(x) < p_{k+1}(x)$, which in turn means that

$$\begin{aligned} \bar{q}_{k+1}(x) &\equiv \bar{\beta}^{-1}(1 - \bar{p}_{k+1}(x)) \geq \beta^{-1}(1 - \bar{p}_{k+1}(x)) \\ &\geq \beta^{-1}(1 - p_{k+1}(x)) \equiv q_{k+1}(x), \end{aligned}$$

and we are done.

Q.E.D.

REFERENCES

1. R. AUMANN, Agreeing to disagree, *Ann. Statist.* 4 (1976), 1236-1239.
2. G. BROCK, "The U.S. Computer Industry: A Study of Market Power," Ballinger, Cambridge, Mass., 1975.
3. A. DIXIT, A model of duopoly suggesting a theory of entry barriers, *Bell J. Econ.* 10 (1979), 20-32.
4. D. EASLEY, R. T. MASSON, AND R. J. REYNOLDS, A dynamic analysis of predatory pricing with rational expectations, mimeo, Cornell University, 1981.

5. J. HARSANYI, Games of incomplete information played by Bayesian players (parts I, II, III), *Manage. Sci.* 14 (1967-68), 159-182, 320-334, 486-503.
6. D. KREPS, P. MILGROM, J. ROBERTS, AND R. WILSON, Rational cooperation in the finitely repeated prisoners' dilemma, *J. Econ. Theory* 27 (1982), 245-252.
7. D. KREPS AND R. WILSON, Reputation and imperfect information, *J. Econ. Theory* 27 (1982), 253-279.
8. D. KREPS AND R. WILSON, Sequential equilibrium, *Econometrica*, in press.
9. M. MACGREGOR, A resolution of the chain-store paradox, draft, University of California, Berkeley, 1979.
10. J. MCGEE, Predatory price cutting: the standard Oil (N.J.) case, *J. Law and Econ.* 1 (1958), 137-169.
11. J. MCGEE, Predatory pricing revisited, *J. Law and Econ.* 23 (1980), 289-330.
12. P. MILGROM, An axiomatic characterization of common knowledge, *Econometrica* 49 (1981), 219-222.
13. P. MILGROM AND J. ROBERTS, Limit pricing and entry under incomplete information: an equilibrium analysis, *Econometrica* 50 (1982), 443-459.
14. R. ROSENTHAL, Games of perfect information, predatory pricing and the chain store paradox, *J. Econ. Theory* 25 (1981), 92-100.
15. F. M. SCHERER, "Industrial Market Structure and Economic Performance," 2nd ed., Rand McNally, Chicago, 1980.
16. R. SEITZ, The chain store paradox, *Theory and Decision* 9 (1978), 127-159.
17. M. SPENCE, Entry, capacity investment and oligopolistic pricing, *Bell J. Econ.* 8 (1979), 534-544.
18. U.S. Department of Justice, "The Regulated Ocean Shipping Industry," Washington, 1977.
19. B. YAMEY, Predatory price cutting: notes and comments, *J. Law and Econ.* 15 (1972), 137-147.