

Reputation and Imperfect Information

DAVID M. KREPS AND ROBERT WILSON

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

Received June 18, 1980; revised June 22, 1981

A common observation in the informal literature of economics (and elsewhere) is that in multistage "games," players may seek early in the game to acquire a reputation for being "tough" or "benevolent" or something else. But this phenomenon is not observed in some formal game-theoretic analyses of finite games, such as Selten's finitely repeated chain-store game or in the finitely repeated prisoners' dilemma. We reexamine Selten's model, adding to it a "small" amount of imperfect (or incomplete) information about players' payoffs, and we find that this addition is sufficient to give rise to the "reputation effect" that one intuitively expects. *Journal of Economic Literature*, Classification Numbers: 026, 213, 611.

1. INTRODUCTION

The purpose of this paper is to present some game-theoretic models that illustrate the role of a firm's reputation. Allusions to reputational effects recur in the industrial organization literature on imperfect competition, but formal models and analyses have been lacking. Scherer [21], for example, points to

the demonstration effect that sharp price cutting in one market can have on the behavior of actual or would-be rivals in other markets. If rivals come to fear from a multimarket seller's actions in Market A that entry or expansion in Markets B and C will be met by sharp price cuts or other rapacious responses, they may be deterred from taking aggressive actions there. Then the conglomerate's expected benefit from predation in Market A will be supplemented by the discounted present value of the competition-inhibiting effects its example has in Markets B and C. (page 338)

The intuitive appeal of this line of reasoning has, however, been called the "chain-store paradox" by Selten [24], who demonstrates that it is not supported in a straightforward game-theoretic model. We shall elaborate Selten's argument later, but the crux is that, in a very simple environment, there is no means by which thoroughly rational strategies in one market could be influenced by behavior in a second, essentially independent market.

What is lacking, apparently, is a plausible mechanism that connects behavior in otherwise independent markets.

We show that imperfect information is one such mechanism. Moreover, the effects of imperfect information can be quite dramatic. If rivals perceive the *slightest* chance that an incumbent firm might enjoy “rapacious responses,” then the incumbent’s optimal strategy is to employ such behavior against its rivals in all, except possibly the last few, in a long string of encounters. For the incumbent, the immediate cost of predation is a worthwhile investment to sustain or enhance its reputation, thereby deterring subsequent challenges.

The two models we present here are variants of the game studied by Selten [24]; several other variations are discussed in Kreps and Wilson [8]. The first model can be interpreted in the context envisioned by Scherer: A multimarket monopolist faces a succession of potential entrants (though in our model the analysis is unchanged if there is a single rival with repeated opportunities to enter). We treat this as a finitely repeated game with the added feature that the entrants are unsure about the monopolist’s payoffs, and we show that there is a unique “sensible” equilibrium where, no matter how small the chance that the monopolist actually benefits from predation, the entrants nearly always avoid challenging the monopolist for fear of the predatory response. The second model enriches this formulation by allowing, in the case of a single entrant with multiple entry opportunities, that also the incumbent is uncertain about the entrant’s payoffs. The equilibrium in this model is analogous to a price war: Since the entrant also has a reputation to protect, both firms may engage in battle. Each employs its aggressive tactic in a classic game of “chicken,” persisting in its attempt to force the other to acquiesce before it would itself give up the fight, even if it is virtually certain (at the outset) that each side will thereby incur short-run losses.

After reviewing Selten’s model in Section 2, we analyze these two models in Sections 3 and 4, respectively. In Section 5 we discuss our results and relate them to some of the relevant literature. In particular, this issue of the *Journal* includes a companion article by Milgrom and Roberts [13] that explores many of the issues studied here in models that are richer in institutional detail. Their paper is highly recommended to the reader.

2. THE CHAIN-STORE PARADOX

The models we analyze are variations on the chain-store game studied by Selten [24]. Consider a sequential game with two players called the *entrant* (or potential entrant) and the *monopolist*. The entrant moves first, electing either to *enter* or to *stay out*. Following entry, the monopolist chooses either to *acquiesce* or to *fight*. If the entrant stays out, the incumbent is not called

		PAYOFFS	
		ENTRANT'S	MONOPOLIST'S
ENTRANT STAYS OUT	ENTRANT STAYS OUT	0	a
	ENTRANT ENTERS	b	0
	ENTRANT ENTERS	$b - 1$	-1

FIG. 1. Selten's chain-store game.

upon to move. Payoffs to the players, depending on the moves selected, are given in Fig. 1. We consider the case that $a > 1$ and $0 < b < 1$.

How will this game be played? If the entrant enters, the monopolist chooses between the payoffs 0 if it acquiesces and -1 if it fights, so surely it will acquiesce. Anticipating this response, the entrant chooses between 0 if it stays out and b if it enters, and so it will enter. This is one Nash equilibrium of the game, but there is another: If the entrant were to anticipate that the monopolist would fight entry, then the entrant would want to stay out. Note that it costs the monopolist nothing to adopt the strategy "fight if entry" if no entry occurs. So this is a second Nash equilibrium. But this second equilibrium is not so plausible as the first. It depends on an expectation by the entrant of the monopolist's behavior that, faced with the *fait accompli* of entry, would be irrational behavior for the monopolist. In the parlance of game theory, the second equilibrium is *imperfect*. We suppose that the entrant adopts the "rational expectation" that the monopolist will acquiesce to entry, and we expect the first equilibrium to ensue.

Consider next the case that the game in Fig. 1 is played a finite number of times. A single monopolist plays a succession of N different entrants, where the monopolist's total payoff is the sum of its payoffs in the N stage games. Allow the later entrants to observe the moves in all earlier stages of the game. Scherer's reasoning predicts that in this case the "reputation" effect might come to life: The monopolist, by fighting any early entry, might convince later opponents that it will fight, thus deterring later entries. Indeed, if this were the case, then also the early round opponents would not enter, not wishing to be abused for demonstration purposes. However, as Selten argues, this does not withstand scrutiny. In the last stage the monopolist will not fight because there are no later entrants to demonstrate for. So in the last stage, entry will surely occur. But then in the penultimate stage, the monopolist again has no reason to fight—it is costly in the short run and has no effect on the last stage. The next-to-last entrant, realizing this, will surely enter. This logic can be repeated, unraveling from the back: In each stage entry and acquiescence will occur. To be precise, this is the unique perfect Nash equilibrium of the game; cf. Selten [22, 23, 24]. Apparently, this model is inadequate to justify Scherer's prediction that reputational effects will play a role.

3. ONE-SIDED UNCERTAINTY

Our contention is that this inadequacy arises because the model does not capture a salient feature of realistic situations. (This contention was made first by Rosenthal [17], whose work we shall discuss in Section 5.) In practical situations, the entrants cannot be *certain* about the payoffs to the monopolist. They may be unsure about the monopolist's costs, or they may be uncertain about nonpecuniary benefits that the monopolist reaps—this may be a monopolist who enjoys being tough. The latter might be more colorfully stated by saying that the monopolist plays tough "irrationally"; according to Scherer [21, p. 247], "... fear of irrational or avowedly rapacious action, then, rather than the expectation of rational pricing responses, may be what deters the potential new entrant from entering on a large scale." For whatever reason, the entrants may initially assess some positive probability p that the monopolist's payoffs are not as in Fig. 1 but rather (in the simplest case) as in Fig. 2, reflecting a short-term benefit from a fighting response. In this case, later entrants, observing earlier moves, will revise their assessment p on the basis of what they see. Perhaps in this case the reputation effect will come alive.

We model this formally as follows. There are $N + 1$ players, for N a positive integer. One of the players is the *monopolist*; the others are called *entrant N* , *entrant $N - 1$* , ..., *entrant 1*. The monopolist plays the game in Fig. 1 against each entrant in turn: First it plays against entrant N , then $N - 1$, etc. (We always index time backwards, and we refer to stage n as that part of the game that involves entrant n .) The payoffs for each entrant are given in Fig. 1.

The monopolist's payoffs are more complex: Its total payoff is the sum (undiscounted for now) of its payoffs in each stage, where the stage payoffs are either *all* as in Fig. 1 or *all* as in Fig. 2. The monopolist knows which payoff structure obtains. The entrants, on the other hand, initially assess probability δ that the monopolist's payoff structure is the second one. As the game progresses, each entrant (and the monopolist) observes all prior moves. Consequently, the history of moves prior to stage n may enable entrant n to revise this assessment if the history reveals some information about the relative likelihoods of the monopolist's two possible payoff structures.




FIG. 2. Payoffs for a tough monopolist.

This model conforms to Harsanyi's [7] formulation of a game with incomplete information. Alternatively, it is a game with imperfect information (among the entrants) and perfect recall, in which "nature" initially determines the monopolist's payoff structure, and nature's move is observed by the monopolist but not the entrants. In line with the first interpretation, we refer to the *weak monopolist* or the *strong monopolist*, meaning the monopolist if its payoffs are as in Fig. 1 or Fig. 2, respectively.

Since the players have perfect recall, there is no loss of generality in restricting attention to behavior strategies (Kuhn [10]). We wish to identify a Nash equilibrium of this game and, moreover, we wish the equilibrium identified to be *perfect*. That is, we wish to exclude equilibria that are based on expectations by one player of another's behavior that would not be rational for the latter to carry out if called upon to do so. Because our games have incomplete information, Selten's [22] concept of subgame perfection is inadequate. His concept of "trembling-hand" perfection (Selten [23]), on the other hand, is difficult to employ in games with strategy spaces as complex as those present here. So we use an analogous equilibrium concept called a *sequential equilibrium*. This is a refinement for extensive games of the usual Nash equilibrium that captures the spirit of Selten's perfectness criterion but that is much easier to apply. General definitions and properties of sequential equilibria are given in Kreps and Wilson [9], which we summarize here.

There are three basic parts to the definition of a sequential equilibrium: (a) Whenever a player must choose an action, that player has some probability assessment over the nodes in its information set, reflecting what that player believes has happened so far. (b) These assessments are consistent with the hypothesized equilibrium strategy. For example, they satisfy Bayes' rule whenever it applies. (c) Starting from *every* information set, the player whose turn it is to move is using a strategy that is optimal for the remainder of the game against the hypothesized future moves of its opponent (given by the strategies) and the assessment of past moves by other players and by "nature" (given by the assessment over nodes in the information set). The difference between this and the standard concept of a Nash equilibrium is that (c) is required for *every* information set, including those that will not be reached if the equilibrium strategies are followed. So each player will be willing to carry out its strategy at *every* point in the game, if ever it is called upon to do so. The properties are: Sequential equilibria exist for all finite extensive games. They are subgame perfect Nash equilibria. For a fixed extensive form and probabilities of nature's moves, as we vary the payoffs it is generic that all strict sequential equilibria are trembling-hand perfect, and the equilibrium path of each sequential equilibrium is an equilibrium path for some trembling-hand perfect equilibrium. Every trembling-hand perfect equilibrium is sequential.

In the context of the game analyzed here, the definition of sequential

equilibrium specializes as follows. An equilibrium comprises a (behavior) strategy for each player *and*, for each stage $n = N, \dots, 1$, a function p_n taking histories of moves up to stage n into numbers in $[0, 1]$ such that: (a) Starting from any point in the game where it is the monopolist's move, the monopolist's strategy is a best response to the entrants' strategies. (b) For each n , entrant n 's strategy (contingent on a history h_n of prior play) is a best response to the monopolist's strategy *given that* the monopolist is strong with probability $p_n(h_n)$. (c) The game begins with $p_N = \delta$. (d) Each p_n is computed from p_{n+1} and the monopolist's strategy using Bayes' rule whenever possible. (We will not write (d) precisely—it will be transparent when we give the equilibrium below. That (d) implies “consistency of beliefs” in the sense of Kreps and Wilson [9] may not be apparent, but it does follow from the simple structure of the game being considered here.) The interpretation is that p_n gives the probability assessed by entrant n that the monopolist is strong as a function of how the game has been played up to stage n . Note that in (a) the monopolist's assessment over nodes in its information set is omitted, because all of its information sets are singletons.

We now give a sequential equilibrium for this game. This particular sequential equilibrium has the fortuitous property that, in terms of play from stage n on, p_n is a sufficient statistic for the history of play up to date n . That is, the choices of the players at stage n depend only on p_n and (for the monopolist) the move of entrant n ; and p_n is a function of p_{n+1} and the moves at stage $n + 1$. We are lucky to be able to find a sequential equilibrium with this simple structure; it is not generally the case that one can find sequential equilibria for which the players' assessments are sufficient statistics for past play. (See remark (A) below.)

We begin by giving the functions p_n . Set $p_N = \delta$. For $n < N$, if the history of play up to stage n includes *any* instance that entry was met by acquiescence, set $p_n = 0$. If every entry so far has been met by fighting, and if k is the smallest index ($> n$) such that there was entry at stage k , then set $p_n = \max(b^{k-1}, \delta)$. If there has been no entry, set $p_n = \delta$.

This corresponds to the following recursive definition:

- (a) If there is no entry at stage $n + 1$, then $p_n = p_{n+1}$.
- (b) If there is entry at stage $n + 1$, this entry is fought, and $p_{n+1} > 0$, then $p_n = \max(b^n, p_{n+1})$.
- (c) If there is entry at stage $n + 1$ and either this entry is met by acquiescence or $p_{n+1} = 0$, then $p_n = 0$.

Now that we have described how p_n is computed at every node in the game tree, we can give the strategies of the players in terms of p_n .

Strategy of the Monopolist

- (a) If the monopolist is strong, it always fights entry.

(b) If the monopolist is weak and entry occurs at stage n , the monopolist's response depends on n and p_n : If $n = 1$, the monopolist acquiesces. If $n > 1$ and $p_n \geq b^{n-1}$, the monopolist fights. If $n > 1$ and $p_n < b^{n-1}$, the monopolist fights with probability $((1 - b^{n-1})p_n)/((1 - p_n)b^{n-1})$ and acquiesces with the complementary probability. (Note that when $p_n = 0$, the probability of fighting is zero, and when $p_n = b^{n-1}$, the probability of fighting is one.)

Strategies of the Entrants

If $p_n > b^n$, entrant n stays out. If $p_n < b^n$, entrant n enters. If $p_n = b^n$, entrant n randomizes, staying out with probability $1/a$.

PROPOSITION 1. *The strategies and beliefs given above constitute a sequential equilibrium.*

Proof. We only sketch the proof, leaving details to the reader. In the context of this game, there are two things to verify: First, the beliefs of the entrants must be consistent with the strategy of the monopolist, in the sense that Bayes' rule holds whenever it applies. Second, starting from any information set in the game, no player has the incentive (in terms of the payoff for the remainder of the game) to change its selection of move at that information set. For entrants, this verification is made using the beliefs given above. (Once this is verified, the Bellman optimality principle together with the fact that beliefs are Bayesian consistent ensures that no player can unilaterally change its strategy and benefit starting from any point in the game tree.)

The verification of Bayesian consistency is easy. If no entry takes place at stage n , nothing is learned about the monopolist, and we have $p_{n-1} = p_n$ in such instances. If $p_n \geq b^{n-1}$, then the monopolist is supposed to fight entry. If $p_n = 0$, then the monopolist is supposed to acquiesce. So in these cases, Bayes' rule implies that $p_{n-1} = p_n$ (as long as the monopolist follows its strategy). In each case, this is what we have. Finally, for $p_n \in (0, b^{n-1})$, there are positive probabilities that the monopolist will acquiesce and that it will fight entry. It only acquiesces if it is weak, and, indeed, in this case we have $p_{n-1} = 0$. If it fights, Bayes' rule requires that

$$\begin{aligned} p_{n-1} &= \text{Prob(monopolist strong} \mid \text{monopolist fights)} \\ &= \text{Prob(monopolist strong and fights)} / \text{Prob(fights)} \\ &= \frac{\text{Prob(fights} \mid \text{strong)} \cdot \text{Prob(strong)}}{\text{Prob(fights} \mid \text{strong)} \cdot \text{Prob(strong)} + \text{Prob(fights} \mid \text{weak)} \cdot P(\text{weak})} \\ &= \frac{1 \cdot p_n}{1 \cdot p_n + [((1 - b^{n-1})p_n)/((1 - p_n)b^{n-1})][1 - p_n]} = b^{n-1}, \end{aligned}$$

which is what we have posited. Thus beliefs and strategies are Bayesian consistent.

Note that there are two instances in which Bayes' rule does not apply: $p_n \geq b^{n-1}$ and the monopolist acquiesces to entry; $p_n = 0$ and the monopolist fights. In each case we set $p_{n-1} = 0$. In words, we assume that any acquiescence is viewed by the entrants as "proof" that the monopolist is weak, and the entrants are unshakeable in this conviction once it is formed. This assignment of beliefs off the equilibrium path is *somewhat* arbitrary—there are other assessments that work as well. But this assignment is not wholly capricious—there are assessments that would not give an equilibrium. (This will be discussed more fully below.)

(Repeating an earlier contention, this set of assessments is consistent in the sense of Kreps and Wilson [9]. A direct proof is not difficult.)

Verification that the entrants are playing optimally is straightforward. If $p_n \geq b^{n-1}$, entrant n expects entry to be fought, and so it stays out. If $p_n \in (b^n, b^{n-1})$, acquiescence will occur with positive probability, but with probability less than $1 - b$. Again it is better to stay out. If $p_n = b^n$, acquiescence follows entry with probability $1 - b$, and the entrant is indifferent. If $p_n < b^n$, the probability of acquiescence exceeds $1 - b$, and the entrant enters.

To see that the strong monopolist is playing optimally, note that if the entrants follow the strategy above, acquiescence at any point results in more future entries than does fighting. In the short run fighting is better for the strong monopolist, and in the long run fewer entries are better, so the strong monopolist will always fight.

Finally, for the weak monopolist, one can verify inductively that given that these strategies are followed from stage n to stage 1, the expected payoff to the weak monopolist from stages n to 1 is given by the following function of p_n :

$$\begin{aligned} v_n(p_n) &= a(t - k(p_n) + 1) + 1 && \text{if } b^n < p_n = b^{k(p_n)-1}, \\ &= a(t - k(p_n) + 1) && \text{if } b^n < p_n < b^{k(p_n)-1}, \\ &= 1 && \text{if } p_n = b^n, \text{ and} \\ &= 0 && \text{if } p_n < b^n, \end{aligned}$$

where $k(p) = \inf\{n: b^n < p\}$ for $p > 0$, and $k(0) = \infty$. Now suppose that entry occurs at stage n . By acquiescing, the monopolist receives zero both in this stage and in the rest of the game (since p_{n-1} will be set equal to zero). By fighting, the monopolist receives -1 in this stage and future expected payoffs of 0 if $p_n = 0$, 1 if $p_n \in (0, b^{n-1}]$, and more than 1 if $p_n > b^{n-1}$. Thus the weak monopolist is happy to follow the strategy given above. ■

It is easiest to understand the nature of this equilibrium by tracing through the play of the game for “typical” values of δ and b , say $\delta = 1/10$ and $b = 1/2$. Note that in this case, $k(\delta) = 4$. Refer to Fig. 3. At stage N (presumed to be greater than 4) the game begins with $p_N = \delta$. At this stage, the monopolist would fight entry regardless of its payoffs, so entry is forestalled. The game evolves along arrow (a) to the point $p_{N-1} = p_N = \delta$. Note that *if* there is entry, the monopolist is willing *ex post* to fight—to acquiesce moves the game along arrow (b) to $p_{N-1} = 0$, from which point the monopolist nets zero. Fighting costs 1 immediately, but acquiescing costs much more in the future. (Note that all that is necessary is that acquiescence cost at least one—as long as acquiescence resulted in $p_{N-1} \leq 1/16$ this would




FIG. 3. Temporal evolution of the game.

be so in this case. So here is one place where the non-Bayesian reassessments need not be *precisely* as specified above to have an equilibrium. But note that $p_{N-1} \leq 1/16$ is necessary—otherwise the weak monopolist would rather acquiesce than fight at this stage.) The game continues in this fashion (arrows (c), (e), and (f)) until date 4 ($=k(\delta)$). At this date the monopolist might acquiesce if it is weak and if it is challenged—the strategy of the weak monopolist is chosen so that following entry, acquiescence leads to $p_3 = 0$ (arrow (g)) and fighting leads to $p_3 = b^3$ (arrow (h)). But this does *not* give entrant 4 enough incentive to enter—the game actually evolves along arrow (i) to $p_3 = \delta$. At date 3 the weak monopolist again will randomize if challenged (so that arrows (j) and (k) give the posteriors), and now there is high enough probability of acquiescence for entrant 3 to enter. If the monopolist acquiesces, the game moves along arrow (j) to $p_2 = 0$. At this point the monopolist is known to be weak, and entrants 2 and 1 both enter with the monopolist acquiescing each time. The monopolist thereafter is supposed to acquiesce; if it fights instead at, say, date 2, entrant 1 disregards this and continues to believe that the monopolist is weak. That is, $p_1 = 0$ if either the monopolist acquiesces *or* if it fights. (Note that we could have $p_1 \leq 1/2$ if the monopolist fights, and still we would have an equilibrium. But if $p_1 > 1/2$, then the weak monopolist would prefer to fight, upsetting the equilibrium. Again there is some freedom in defining beliefs off the equilibrium path, but not complete freedom.) Back at stage 3, if the monopolist fights entry, the game moves along arrow (k) to $p_2 = 1/4$. At this point entrant 2 is indifferent between entering and staying out, and chooses between the two randomly. If entrant 2 enters, the weak monopolist randomizes between acquiescence (arrow (n)) and fighting (arrow (p)). If entrant 2 stays out, the game moves along arrow (q) to $p_1 = 1/4$ (and entrant 1 surely enters).

The remarkable fact about this equilibrium is that even for very small δ , the “reputation” effect soon predominates. Even if the entrants assess a one-in-one-thousand chance that the monopolist would prefer (in the short run) to fight, if there are more than ten stages to go the entrant stays out because the monopolist will *surely* fight to preserve its reputation. Note the “discontinuity” that this causes as the number of stages in the game goes to infinity:

$$\begin{aligned} \lim_{N \rightarrow \infty} v_N(\delta)/N &= a && \text{if } \delta > 0, \text{ and} \\ &= 0 && \text{if } \delta = 0. \end{aligned}$$

The obvious question at this point is: To what extent is this equilibrium unique? It is not the case that it is the unique Nash equilibrium for this game, for the following four reasons.

- (a) There are other Nash equilibria that are not sequential equilibria.

(That is, that are not, roughly speaking, perfect.) For example, it is a Nash equilibrium for the monopolist to fight any entry (regardless of its payoffs), and for the entrants never to enter. But this behavior is not “ex post rational” for the weak monopolist in stage one. In general, we wish to allow only sequential equilibria, and we confine attention to those for the remainder of this discussion.

(b) There are sequential equilibria where the strong monopolist acquiesces to entry. For example, if $N = 2$, $b = 1/2$, and $\delta = 2/3$ (very high probability that the monopolist is strong), it is a sequential equilibrium for entrant 2 to enter, the monopolist to acquiesce to this entry regardless of its payoffs, and for entrant 1 to adopt the strategy: Stay out if the monopolist acquiesces in stage 2; enter if the monopolist fights in stage 2. (In stage one, the monopolist responds with its ex post dominant action.) This is sequential because it is supported by the following beliefs of entrant 1:

$$\text{Prob(monopolist strong | acquiescence in stage 2)} = p_2 = 2/3,$$

$$\text{Prob(monopolist strong | fight in stage 2)} = 1/4.$$

For the given strategies, the first of these reassessments follows from Bayes’ rule, and the second is “legitimate” because Bayes’ rule does not apply: There is zero prior probability that the monopolist will fight in stage 2.

Although this is a sequential equilibrium, we contend that it is not very sensible. The flaw is in the beliefs of entrant one—if there is fighting in stage 2, entrant 1 revises *downward* the probability that the monopolist is strong. Intuitively it seems at least as likely that the strong monopolist would defect and fight as that the weak monopolist would do so. Thus it seems intuitive that entrant one will assess

$$\text{Prob(strong | fight)} \geq \text{Prob(strong | acquiesce)}.$$

But if we insist on this condition holding, then the equilibrium given immediately above is excluded.

Putting this formally, we will call the beliefs $\{p_n\}$ of the entrants *plausible* if given two histories h_n and h'_n of play up to stage n , if h_n and h'_n are the same except that some plays of “fight” in h_n are “acquiesce” in h'_n , then $p_n(h_n) \geq p_n(h'_n)$. We wish to allow only sequential equilibria that are supported by plausible beliefs. Note that this is not true of the equilibrium immediately above, but it is true of the equilibrium Proposition 1.

(c) In the sequential equilibrium given in this section, there is some freedom in describing what happens off the equilibrium path. For example, we have said that if $p_n = 0$ and the monopolist fights entry, then the entrants set $p_{n-1} = 0$. Thus once $p_n = 0$ in our equilibrium, every subsequent entrant enters. But we would also have an equilibrium if we set $p_{n-1} = b^{n-1}$ after

such a defection from the equilibrium, and then entrant $n - 1$ would randomize between entering and staying out. Note well, this concerns the behavior of entrant $n - 1$ only off the equilibrium path, but in terms of strategies it is a different equilibrium. We cannot hope to have uniqueness off the equilibrium path.

(d) Finally, there is a bit of freedom in defining equilibria along the equilibrium path when $\delta = b^n$ for some $n \leq N$: The behavior of entrant n in this case need not conform to the strategy above—any randomization will work.

Except for these four problems, we do get uniqueness:

PROPOSITION 2. *If $\delta \neq b^n$ for $n \leq N$, then every sequential equilibrium with plausible beliefs has on-the-equilibrium-path strategies as described previously. Thus every sequential equilibrium with plausible beliefs has the value functions given above.*

The proof is by induction and is left to the reader. We simply note that in carrying out the induction one establishes the following:

(a) The value function of the strong monopolist (in equilibrium) will be a nondecreasing function of p_n , and the strong monopolist will therefore fight any entry.

(b) The value function of the weak monopolist will be a nondecreasing function of p_n and will be given by the formula in the proof of Proposition 1 for $\delta \neq b^n$, $m \leq n$.

(c) If there is entry at stage n and if the monopolist fights this entry, then entrant $n - 1$ must stay out with probability exactly $1/a$.

By going through this proof, the reader will see the intuition behind this equilibrium, which we will try to summarize here. As long as beliefs are plausible, the strong monopolist will always fight entry. Thus any acquiescence is conclusive proof that the monopolist is weak. Moreover, such evidence once given must result in zero payoff for the monopolist—the argument of Selten that we have given in Section 1 applies (with minor modifications). If entrant n is to enter, then it must be that there is probability $1 - b$ (at least) that the monopolist will acquiesce, which requires that the weak monopolist is randomizing or simply acquiescing. This also requires that $p_n \leq b$, and, from Bayes' rule, that if this entry is met by fighting, then $p_{n-1} \geq p_n/b$. Thus if we begin with $\delta > b^m$, there can be at most m entrants who have a positive probability of entering. As N gets large, then, the value to either monopolist must asymptote to aN , and, for $N > 2m$, the weak monopolist would always wish to fight entry. (In fact, in the equilibrium it turns out that this is true for $N \geq m$.)

We close this section by listing several extensions and embellishments of the basic model.

(A) We have dealt above with the case $a > 1$. If $0 < a \leq 1$, then the same basic structure for the equilibrium emerges, in that for sufficiently large n , entrants do not enter because the monopolist will fight with probability one. The play near the end of the game is more complicated however. In particular, one cannot obtain an equilibrium where entrant n 's strategy depends only on p_n —it depends instead on p_n and the history of play in the last j rounds, where j is the smallest integer such that $ja > 1$.

(B) If the monopolist discounts its payoffs by a factor ρ per period, the following results. If $\rho > 1/a$, then the equilibrium is *precisely* as above except that the randomizing probabilities of the entrants must change. If $\rho \leq 1/(a + 1)$, then the equilibrium is quite different—the weak monopolist acquiesces at the first entry, so entrants enter if $p_n < b$ and stay out if $p_n > b$. For ρ such that $1/(a + 1) < \rho \leq 1/a$, the basic character of the equilibrium is just as in the case of $\rho > 1/a$ —for large enough n entrants stay out because the monopolist will fight any entry. But the equilibrium is complicated for small n , resembling the equilibrium in the undiscounted case where $a < 1$.

(C) Suppose that instead of the sequential game depicted in Fig. 1, each stage consists of a two-by-two simultaneous move game, Table I, where the payoffs with probability $1 - \delta$ are shown in (a) and the payoffs with probability δ are shown in (b). (We assume $0 < b < 1$ and $a > 1$.) Otherwise the structure of the game is the same: One of these two bimatrices is chosen at the outset, according to the probabilities given. One monopolist plays against N entrants in sequence. The monopolist knows which bimatrix was chosen; the entrants do not.

For $\delta = 0$, the argument of Selten is easily adapted to show that the *unique* equilibrium (perfect or not) has row 2, column 1 played in each stage. This is because row 2 is strongly dominant in the stage game. But for $\delta > 0$ we get an equilibrium almost identical to the one discussed above: For stages n such that $b^n < \delta$, the monopolist plays row 1 regardless of which bimatrix was selected, and the entrant responds with column 2. (The play of the game is a bit different near the end of the game.) So we see that a little incomplete

TABLE I

		entrant	
		column 1	column 2
monopolist	row 1	-1, b-1	a-1, 0
	row 2	0, b	a, 0

(a)

		entrant	
		column 1	column 2
monopolist	row 1	0, b-1	a, 0
	row 2	-1, b	a-1, 0

(b)

information can not only make an imperfect equilibrium perfect (more accurately, sequential)—it can also make as part of a sequential equilibrium the play of an action that with very high probability is *strongly dominated* in the stage game.

(D) Paul Milgrom has pointed out to us that similar equilibria can be found even when every player in the game knows the payoffs of the monopolist, as long as this knowledge is not *common knowledge*. That is suppose all the entrants know the monopolist's payoffs, but they are not certain whether their fellow entrants have this information. Then (with the proper precise specification) they fear that the weak monopolist will fight (for large n), in order to maintain its reputation among the other entrants. This being so, the entrant will not enter. And the monopolist, even if it knows that all the entrants know that it is weak, may be willing to fight entry early on, in order to help "convince" subsequent entrants that it (the monopolist) is not sure that the entrants know this. (Precise arguments of this form are found in Milgrom and Roberts [13].) Selten's argument requires that it is common knowledge that the monopolist is weak. In real-life contexts this is a very strong assumption, and weakening it ever so slightly (more slightly than we have done above) can give life to the "reputation" effect.

(E) We have dealt exclusively with the case of a single monopolist playing against N different entrants. Another interesting case is where a single monopolist plays N times against a single entrant. For the game we have analyzed in this section, this turns out to have no effect on the equilibrium. (We leave this to the reader to verify.) But as we shall see in the next section, this is due (at least in part) to the fact that there is no uncertainty about the payoffs of the entrants.

4. TWO-SIDED UNCERTAINTY

In this section we consider what happens when the monopolist is unsure about the payoffs of the entrants. The most interesting formulation of this problem is where a single monopolist plays the stage game of Fig. 1 a total of N times against a single opponent. The payoff to each player is the sum of the player's payoffs in each stage. The monopolist's payoffs are as in Fig. 1 or Fig. 2, with probabilities $1 - \delta$ and δ , respectively. The entrant's payoffs are as in Fig. 1, for some b such that $0 < b < 1$ with probability $1 - \gamma$ and for some other $b > 1$ with probability γ . Each player knows its own payoffs at the start of the game, and each is unsure of the payoffs of its opponent. The payoffs are statistically independent.

Continuing the terminology of Section 3, we shall refer to the weak entrant

as the entrant if its payoffs satisfy $0 < b < 1$, and the strong entrant if its payoffs satisfy $b > 1$.

Note that the strong entrant does better to enter than to stay out in any stage, even if the monopolist is sure to fight. Because it seems plausible that entry will not decrease the probability that the monopolist will acquiesce subsequently, we look for equilibria where the strong entrant always enters. Thus any failure to enter brands the entrant as weak, at which point we are back to the situation of Section 3. (Recall that it did not matter there whether there was a single entrant or N entrants.) Similarly, we look for an equilibrium where the strong monopolist always fights. Thus any failure to fight brands the monopolist as weak, following which the entrant always enters and the monopolist always acquiesces. We search, then, for an equilibrium of the following sort: The strong entrant always enters. The strong monopolist always fights. The weak entrant chooses a strategy that is a mixture of "stopping rules": A stopping rule gives the date at which the entrant will "give in" and not enter if the monopolist has not acquiesced yet. (The entrant may later re-enter, as we will then follow the equilibrium of Section 3.) The weak monopolist will also mix among stopping rules: A stopping rule for the monopolist gives the date at which the monopolist will first acquiesce if the entrant has not retreated first. If one side or the other gives in, we move to either the situation of Section 3 or to where entry-acquiescence follows until the game ends.

Giving a complete specification of the equilibrium that is obtained is extraordinarily tedious, because it is based on some very involved recursions. Still, we can give a rough description of what happens. At any stage n the previous play of the game is summarized into two statistics: p_n , the probability assessed by the entrant that the monopolist is strong; q_n , the probability assessed by the monopolist that the entrant is strong. (The game begins with $p_N = \delta$ and $q_N = \gamma$.) Thus the "state space" of the game at stage n is the unit square, as depicted in Fig. 4. The edge $q_n = 0$ is the subject of




FIG. 4. State space of the game at stage n .

Section 3. The edge $p_n = 0$ can be analyzed using the argument of Selten with the conclusion: The entrant always enters, and the weak monopolist always acquiesces.

The square is divided into two regions by a curve, as shown. If (p_n, q_n) lies in region I, say at the point x , then the entrant enters regardless of its payoffs, and the weak monopolist randomizes. If the weak monopolist acquiesces, the game evolves to the point y (actually, to this point in the *next* square—the square for stage $n - 1$). If it fights, or if it is strong and hence fights, the entrant uses Bayes' rule to compute p_{n-1} , landing at the point z on the curve, and just *beyond* the curve in region II of the next square. If (p_n, q_n) lies in region II, say at x' , then the weak entrant randomizes. If it stays out, the equilibrium of Section 3 ensues—and the next stage begins at the point y' . If it enters (or if it is strong and therefore enters), the monopolist recomputes the probability that it is strong, landing at the point x'' along the curve. Then the monopolist randomizes (if weak), and the game evolves to y'' or to z'' . Both z and z'' are in region II of the next square, so the next round begins with randomization by the weak entrant, and so on.

Except for the very start of the game, when $p_N = \delta$ and $q_n = \gamma$, most of the play takes place along the curve. (Actually, the curve shifts slightly as n changes.) So we see an initial jump to the curve (or to one of the two edges), followed by a slow climb up the curve with ever present the chance that a jump to one of the edges will occur. With probability one, a weak player will eventually give in, so we either jump to an edge eventually or, if both players are strong, we reach the point $p_0 = q_0 = 1$. This is a game of "chicken," where once begun, each side (if weak) randomizes between a small probability of giving in and a large probability of daring the other side for one more round. The relative size of these probabilities is required by the conditions of an equilibrium: Daring once more costs something this round, but giving in is costly for the rest of the game. So it must be that daring once more does not give either player a substantial chance of *immediate* gain; the opponent must be about to "dare" once again with large probability.

While it is tedious to give the exact equilibrium in the discrete time formulation, it is relatively simple to do so in a continuous time version of the game. So we shall now develop that continuous-time version. (We should forewarn the reader: We will be somewhat sloppy in what follows. But everything we say can be made exact.)

To begin, consider the game of Section 3 played against a single entrant over the time period N to 0. Instead of playing at times $N, N - 1, \dots, 1$, for the stakes (per stage) given in Section 3, we imagine that an integer K is given, and that the game is played "more frequently, for reduced stakes," with play at times $N, (KN - 1)/K, (KN - 2)/K, \dots, 1/K$, for stakes $1/K$ times the stakes given. It is the *number* of times that the monopolist has left to demonstrate its "toughness" that is decisive in Section 3, so we find that if $k(\delta) = n$, the

entrant stays out (and the monopolist would surely fight) at all times $t > n/K$. As K goes to infinity, we see that the entrant stays out until “the very last instant” of time.

With this limiting result as motivation, we now consider the “continuous-time” version of the above game, played over a time interval T to 0. At each time $t \in [T, 0]$, the entrant chooses whether to enter or to stay out, and the monopolist chooses whether to fight entry or to acquiesce. A realization of the entrant’s strategy is formalized as a (measurable) function $e: [T, 0] \rightarrow \{0, 1\}$, where $e(t) = 1$ means that the entrant is entering at date t . A realization of the monopolist’s strategy is formalized as a function $f: [T, 0] \rightarrow \{0, 1\}$, where $f(t) = 1$ means that the monopolist is fighting at date t . (We have a “closed-loop” game, so pure strategies would be a pair of functions \mathbf{e} and \mathbf{f} where $\mathbf{e}(t)$ is $F(\mathbf{e}(s), \mathbf{f}(s)), s < t$ -measurable, and $\mathbf{f}(t)$ is $F(\mathbf{f}(s), s < t; \mathbf{e}(s), s \leq t)$ -measurable. We shall not try to be more precise about this here; instead we trust the reader’s ability to see how to formalize what follows.) Given realizations e and f , payoffs to each side are determined by measuring the lengths of times during which there is not entry ($e(t) = 0$), during which entry is fought ($e(t) = 1, f(t) = 1$), and during which there is acquiescence to entry ($e(t) = 1, f(t) = 0$), and assigning payoffs accordingly. For example, if λ denotes Lebesgue measure, then the weak monopolist’s payoff is

$$\lambda\{e(t) = 0\} \cdot a - \lambda\{e(t) = 1, f(t) = 1\}.$$

In this game, an equilibrium calls for the entrant to stay out as long as the monopolist does not acquiesce and to always enter after any acquiescence is observed; for the strong monopolist to fight any entry; and for the weak monopolist to fight as long as it has not acquiesced yet and to acquiesce forever after an acquiescence. The reader can easily verify that this is an equilibrium. Moreover, if by some “mistake” the entrant entered before time 0, the weak monopolist would want to fight: By acquiescing it saves an “instantaneous” one unit, but then it invites entry for the remainder of the game—a substantial loss that outweighs the instantaneous savings.

(The reader is entitled to be somewhat skeptical about this. By moving to a continuous-time formulation, we have obtained some of the features of the supergame (infinitely repeated) formulation. For example, the equilibrium above is “perfect” even if $\delta = 0$, just as in the supergame with $\delta = 0$. But in the case $\delta = 0$, this equilibrium is not the limit of discrete-but-more-rapid equilibrium play. What justifies this particular equilibrium in the case $\delta > 0$ is that it is the limit of discrete equilibria. We shall return to this point after we discuss the case of two-sided uncertainty.)

Now consider the continuous-time game where there is uncertainty on both sides. The formulation is as above, but now there is uncertainty (at the outset) about the entrant’s payoffs as a function of the realizations of e and f .

We are looking for an equilibrium with the following characteristics: (1) The strong monopolist always fights. (2) The strong entrant always enters. (3) By virtue of (1), if the monopolist ever declines to fight an entry, it is revealed as weak. Thereafter, the entrant always enters and the weak monopolist always acquiesces. (4) By virtue of (2), if the entrant ever fails to enter it is revealed as weak. Assuming that the monopolist has not previously been revealed as weak, the game proceeds as above, with the weak entrant staying out until the end and the monopolist always ready to fight.

Just as in the discrete time formulation, an equilibrium with these features can be recast as an equilibrium in "stopping rules" for the weak entrant and monopolist—each choosing the date at which it will "give in" if its opponent has not given in yet. If the entrant gives in first (at date t), then regime (4) above takes effect, with the weak monopolist obtaining at for the rest of the game, and the weak entrant receiving 0. If the monopolist gives in first at t , regime (3) ensues, with the weak entrant receiving bt and the weak monopolist receiving 0. Until one side or the other gives in, the weak monopolist receives -1 per unit of time, and the weak entrant receives $b - 1$ per unit of time. The equilibrium condition is that each player's stopping time should be optimal given the probability distribution of the other's, and given the assumption that the other player, if strong, will never give in. This game is very similar to the "war of attrition" game; cf. Riley [16] and Milgrom and Weber [14]. It is formally equivalent to a two-person competitive auction, where the stopping times are reinterpreted as bids. This observation will be especially useful later when we discuss the connection between this continuous-time formulation and the discrete-time formulation of the game. (We are indebted to Paul Milgrom for acquainting us with the "war of attrition" and for pointing out the relevance of his work with Weber.)

It is easiest and most illustrative to present the equilibrium using a diagram similar to Fig. 4. In Fig. 5 we have the "state space" of this game—the unit square, interpreted exactly as in Fig. 4. The bottom boundary is where the entrant is known to be weak. Along this boundary (excluding the left hand endpoint) the weak monopolist's payoff function (at date t) is $v_t(p, 0) = at$ and the weak entrant's is $u_t(p, 0) = 0$. The left hand boundary is where the monopolist is known to be weak—here (including the bottom endpoint) $v_t(0, q) = 0$ and $u_t(0, q) = bt$.

The nature of the equilibrium is just as in the discrete case: The state space is divided into two regions by a curve $f(p, q) = 0$ that passes through the points $(0, 0)$ and $(1, 1)$. If the initial data of the game place us in region I, then the game begins with the entrant entering for sure and the monopolist (if weak) randomizing between fighting and immediate capitulation. This randomization is such that if the monopolist does fight at time T , the entrant revises its assessment that the monopolist is strong so as

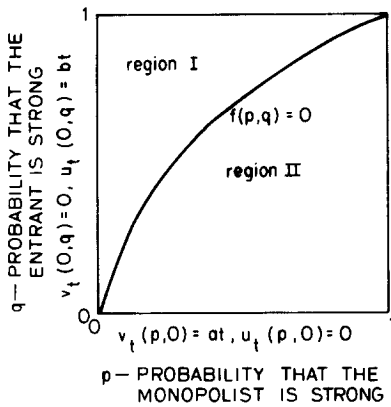


FIG. 5. State space of the continuous-time game.

to go to the curve $f(p, q) = 0$. From region II it is the (weak) entrant that randomizes between immediate capitulation and entry—if it does enter, the monopolist revises its assessment that the entrant is strong up to a point where the curve $f(p, q)$ is reached. Thereafter, the weak monopolist and weak entrant randomize “continuously” between keeping up the fight and capitulating—this is done in a fashion so that as long as they continue to fight, the Bayesian reassessments of each side that the other is strong causes (p_t, q_t) to slide up along the curve toward $(1, 1)$. (Of course, if one side or the other capitulates, transition is made to the appropriate boundary.) There is a time $T^0 > 0$ such that by this time, one side or the other (if weak) has given in with probability one—if both sides are strong, at this time the point $(1, 1)$ has been reached, and we remain there until time $t = 0$.

The difference between this equilibrium and the one for the discrete time game (and the reason that this one is so much easier to compute) is that the curve $f(p, q) = 0$ does not change with t in the continuous-time case. This is so because in the continuous-time version of the game, a game of duration $T/2$ is strategically equivalent to a game of duration T , so long as the priors (δ, γ) are the same. All that changes is that everything takes place twice as rapidly—we could as well think of the game taking place at the same speed but for half the stakes. The values are half as large, but nothing else changes.

We now present a heuristic derivation of the equilibrium, assuming that it has the form outlined above. Note first that along the curve the value functions for each side must be identically zero. This is so because (by hypothesis) both sides are randomizing continuously, and one outcome of these randomizations transfers them to points (the lower boundary for the entrant and the left-hand boundary for the monopolist) where the value function is zero. Let $\pi_t(p_t, q_t)$ and $\rho_t(p_t, q_t)$ be the hazard rate functions associated with the weak monopolist’s and entrant’s randomizations at time t

with posteriors p_t and q_t lying along the curve. That is, in the time interval $(t, t - h)$ there is (up to terms of order $o(h)$) probability $(1 - p_t) \cdot \pi_t(p_t, q_t) \cdot h$ that the monopolist will give in, and $(1 - q_t) \cdot \rho_t(p_t, q_t) \cdot h$ that the entrant will give in. Assuming sufficient continuity, if the value functions are to be constant (and zero) along the curve, it is necessary that the expected change in value to the weak monopolist be zero. Up to terms of order $o(h)$, this is

$$-h + [(1 - q_t) \cdot \rho_t(p_t, q_t) \cdot h][at] = o(h).$$

That is, the weak monopolist's immediate cost $-h$ of keeping up the bluff should be offset by the small chance $(1 - q_t) \cdot \rho_t(p_t, q_t) \cdot h$ that the entrant will give in times the large gain at that will accrue in this event. The analogous argument for the entrant gives

$$(b - 1)h + (1 - p_t) \cdot \pi_t(p_t, q_t) \cdot hbt = o(h).$$

Dividing by h and passing to the limit, we get

$$\rho_t(p_t, q_t) = 1/(at(1 - q_t)) \quad \text{and} \quad \pi_t(p_t, q_t) = (1 - b)/(bt(1 - p_t)).$$

Consider next the evolution of the posteriors p_t and q_t . The probability table that the monopolist would construct at date t for the joint probability distribution that the entrant is weak or not and will give in or not in the interval $(t, t - h)$ is given in Table II (up to terms of order $o(h)$). Thus the conditional probability that the entrant is strong, conditional on not giving in over the interval $(t, t - h)$, is

$$q_{t-h} = q_t/[1 - (1 - q_t)\rho_t(p_t, q_t)h].$$

Thus, ignoring terms of order $O(h)$,

$$(q_t - q_{t-h})/h = q_t(1 - q_t)\rho_t(p_t, q_t).$$

Passing to the limit, this gives $\dot{q}_t = (q_t - q_t^2)\rho_t(p_t, q_t) = q_t/(at)$. Similarly $\dot{p}_t = p_t(1 - b)/(bt)$. Thus along the curve we must have

$$dq_t/dp_t = (q_t b)/((1 - b)ap_t).$$

TABLE II

	weak	strong	
give in	$(1 - q_t)\rho_t(p_t, q_t)h$	0	$(1 - q_t)\rho_t(p_t, q_t)h$
not	$(1 - q_t)(1 - \rho_t(p_t, q_t)h)$	q_t	$1 - (1 - q_t)\rho_t(p_t, q_t)h$
	$(1 - q_t)$	q_t	

This is independent of t , and it is easily integrated to give $q_t = k(p_t)^c$ where k is a constant of integration and $c = b/((1 - b)a)$. To ensure that $(1, 1)$ is on the curve, we must have $k = 1$. Therefore the curve is given by

$$f(p, q) \equiv q^{(1-b)/b} - p^{1/a} = 0.$$

(Note well the normalization of k so that $(1, 1)$ is on the curve. This will be important later on.)

We can solve similarly for q_t and p_t . Integrating $\dot{q}_t = q_t/(at)$ yields $q_t = k' t^{-1/a}$. Analogously, $p_t = k'' t^{-(1-b)/b}$. The constants k' and k'' are determined by the initial conditions. Suppose, for example, that we initially have a prior (δ, γ) that lies in region II. Then the initial randomization is by the weak entrant and yields posterior $q_T = \delta^c$ if the entrant does enter. Solving for k' yields $k' = \delta^c T^{1/a}$. Solving for k'' yields $k'' = \delta T^{(1-b)/b}$. Thus

$$p_t = \delta(T/t)^{(1-b)/b} \quad \text{and} \quad q_t = \delta^c(T/t)^{1/a}.$$

Note that these yield $p_t = 1$ and $q_t = 1$ for $t = T\delta^{ac} = T\delta^{b/(1-b)}$. (Of course, both p_t and q_t hit one simultaneously as the curve has been normalized to pass through $(1, 1)$.) The point to note here is that in this equilibrium, the posterior $(1, 1)$ will be reached at a time T° strictly between T and 0 (unless $\delta = 0$ or 1) so long as neither player gives in previously. But of course, the posterior $(1, 1)$ can only be reached if with probability one both of the weak players would have given in. So, according to this equilibrium, if the two players are both strong, they will learn this before the game terminates. Put another way, the date T° previously referred to is $\delta^{b/(1-b)}T$ (for (δ, γ) in region II). The formulae change somewhat for (δ, γ) in region I, but the qualitative conclusions are the same.

Does this heuristic derivation hold up? That is, do we really have an equilibrium? There are two things to worry about. First, in several places the heuristic arguments that we give depend somewhat on sufficient regularity of the functions π_t and ρ_t . The reader can make these arguments rigorous for the functions that we derived. Second (and more substantially), the necessary conditions that were developed for π_t and ρ_t were necessary for the value functions to be *constant* along the curve. To have an equilibrium we require somewhat more: The value functions must be identically *zero* along this curve. This is where the normalization of the curve comes in: At the point $(1, 1)$, the value functions are clearly zero for each weak player, as each is certain that the other is strong. Put another way, suppose (once the curve is reached) that one side or the other is weak and decides *not* to randomize but simply to wait out its opponent. The conditions that gave us π_t and ρ_t ensure that the change in expected value is zero as this goes on. (For the technically minded, apply Dynkin's formula to the appropriate generalized Poisson process.) And at date T° , if nothing has happened the player that is waiting

knows that its opponent is strong—it should immediately give in for a value of zero. Thus the value all along the curve is zero. This, with a little careful argument, suffices to show that we really do have an equilibrium.

Two final comments about this equilibrium seem in order. First, the value functions to each (weak) player are easily computed. In region II, they are $u_T(\delta, \gamma) = 0$ and $v_T(\delta, \gamma) = [(\delta^c - \gamma)/\delta^c] aT$; in region I, they are $u_T(\delta, \gamma) = [(\gamma^{1/c} - \delta)/\gamma^{1/c}] bT$ and $v_T(\delta, \gamma) = 0$. That is, they are simple linear interpolates of the value of zero along the curve and the values along the bottom boundary in region II and the left-hand boundary in region I.

Second, we noted earlier that the continuous-time formulation can introduce equilibria that are not limits of the equilibria for discrete-time models. We should like to know that the continuous-time equilibrium just presented is indeed the limit of the discrete-time equilibrium with which we began this section. We have not checked all the details, but we are quite sure that this is so. To see this, recall that the discrete time game can be posed as an optimal stopping problem where the entrant is limited to stopping at, say, discrete times $T, (TK - 1)/K, \dots, 1/K$ and the monopolist is limited to stopping at times (say) $T - 1/(2K), (TK - 3/2)/K, \dots, 1/(2K)$. The continuous-time problem is one where stopping at any time $t \in [T, 0]$ is possible. It is easy to move from the discrete-to-continuous-time versions of the problem when there is one-sided uncertainty, so we know that we have convergence of the value of “stopping” at particular times. As K goes to infinity, the sets of available strategies also converge, and the methods of Milgrom and Weber [14] apply to show convergence of the equilibria of the discrete games to the continuous-time version. (Indeed, Paul Milgrom has shown us how, by viewing the continuous-time game as a game in distributional strategies, it is simple to derive the equilibrium given above.)

Before concluding this section, we also note that this gives just one sort of formulation of the problem with two-sided uncertainty. We might consider what happens when a single monopolist plays against a succession of different entrants (each of whom plays the game once), where the monopolist is uncertain of the entrants’ payoffs. In such a game we would have to specify the way in which the entrants’ payoffs are related—they might all be identical (perfectly correlated, from the point of view of the monopolist), or they might be independently and identically distributed, or something between these two extremes. Both of these extreme cases are analyzed in detail in Kreps and Wilson [8]. The case of identical entrants gives the most interesting comparison with the model analyzed above: With identical entrants who only play the game once, the first entrant nearly always “tells the truth” by refusing to enter if weak, and the weak monopolist will with substantial probability fight the first few entries, just to keep the weak entrants “honest.” What this illustrates is that the game of “chicken” that we see above requires *both* two-sided uncertainty *and* that each side has a stake

in maintaining its reputation. When it is only one side that will participate in many stages, the other has little motivation to dissemble and will not fight too hard to attain/maintain its reputation. (Another interesting formulation of the problem is where there is a population of entrants and a population of monopolists, and in each round there is a random assignment of one monopolist to one entrant, in the manner of Rosenthal [18] and Rosenthal and Landau [19]. We have done no analysis of this formulation.)

5. DISCUSSION

We have presented these simple examples to illustrate formally the power of “reputation” in finitely repeated games. That reputation is powerful in reality is very well appreciated: In the context of Industrial Organization, recall the quotation from Scherer in Section 1. Consider the importance of reputation in contract and labor negotiations; in a firm’s employment practices; in a firm’s “good name” for its product; in the maintenance of a cartel (or in the prisoners’ dilemma game); in international diplomacy. To each of these contexts, our analytical structure can be applied to yield the conclusions: If the situation is repeated, so that it is worthwhile to maintain or acquire a reputation, and if there is some uncertainty about the motivations of one or more of the players, then that uncertainty can substantially affect the play of the game. There need not be much uncertainty for this to happen. The power of the reputation effect depends on the nature of one’s opponents; notably on whether they also seek to acquire a reputation.

Phenomena that bear the interpretation of “reputation” are not entirely new to the literature of formal game theory. They are implicit in much of the literature on super-games, where the stage game is repeated infinitely often, or where there is always high enough probability of at least one more repetition (Rubinstein [20] is a representative citation). Indeed, Dybvig and Spatt [6] make explicit use of the reputation interpretation in a super-game context. What is new in this paper (and in Milgrom and Roberts [13]) is the observation that with a very little imperfect information, these effects come alive in finitely repeated games. Comparing the two approaches is difficult, but it is worth noting that in the models reported here, the problem of multiplicity of equilibria that plagues the super-game literature is substantially alleviated. Also, we believe that we have interesting models of the sorts of “wars” that might go on between players to see which equilibrium will ensue. But we are far from ready to make a very informed comparison of the two approaches—at this point, we can only claim that this seems to be an interesting alternative way to produce reputation effects.

A point made briefly in Section 3 is worth repeating here. To keep matters simple, we posited the simplest type of uncertainty: Players are uncertain

about the payoffs of their fellows. But it does nearly as well if there is no uncertainty about players' payoffs, but there is uncertainty about whether this is so. In the parlance of game theory, for these effects to disappear, payoffs must be common knowledge. (Milgrom and Roberts [13] present formal models to back up this contention.) This is a very strong assumption for any real-life application.

The reader may object that in order to obtain the reputation effect, we have loaded the deck. That is, we have a model where reputation is easily shattered, making it all the more valuable; there are at most two types of each player; and each player has only two possible actions. To the first of these criticisms we plead guilty: The power of reputation seems to be positively related to its fragility. As for the second, the models of Milgrom and Roberts [13] have continua of types of monopolists, so this does not seem crucial to our conclusions. And to the third, we do admit that this has made it easy for us to get a "pooling equilibrium" (to borrow a term from the insurance literature), where one type successfully mimics another. The analysis of Milgrom and Roberts [12] shows that with a continuum of actions, one can also get screening equilibria in these sorts of models. But this is not necessary: Crawford and Sobel [3] investigate a class of models with a continua of actions where some pooling is necessary in any equilibrium. The assumption of only two actions makes things easier for us, but we doubt that it is crucial.

What is evident from our simple examples is that a very little uncertainty "destabilizes" game-theoretic analysis in games with a fairly large number of stages. The reader may suspect that something more is true: By cleverly choosing the nature of that small uncertainty (precisely—its support), one can get out of a game-theoretic analysis whatever one wishes. We have no formal proposition of this sort to present at this time, but we certainly share these suspicions. If this is so, then the game-theoretic analysis of this type of game comes down eventually to how one picks the initial incomplete information. And nothing in the *theory* of games will help one to do this.

This reinforces a point made by Rosenthal [17]. Rosenthal investigates the original chain-store game and makes the point with which we began: The paradoxical result in Selten's analysis is due to the complete and perfect information formulation that Selten uses. In a more realistic formulation of the game, the intuitive outcome will be predicted by the game-theoretic analysis. Rosenthal does not provide this analysis, despairing of the analyst's ability to solve an adequate formulation. Instead, he suggests an analysis using the paradigm of Decision Analysis, where one tries to assess *directly* how the entrants will respond to early round fighting by the monopolist. Such an analysis can certainly lead to the intuitive outcome, as shown by Macgregor [11]. But, as Rosenthal notes, the weakness in this approach is the *ad hoc* assessment of entrants' behavior. We have carried out a game-

theoretic analysis of *one* very simple incomplete information formulation. We therefore have avoided *ad hoc* assumptions about the entrants' behavior. But we have made *ad hoc* assumptions about their information, and we have found that small changes in those assumptions greatly influence the play of the game. So at some level, analysis of this sort of situation may require *ad hoc* assumptions.

We have presented models in this paper that demonstrate the reputation effect as simply and as powerfully as possible. In order to do this, we have not tried to model realistic settings from Industrial Organization or some other economic context. (Milgrom and Roberts [13] rectify this deficiency: They concentrate somewhat more on the application of these ideas.) To illustrate how these ideas might be applied, we close with two examples.

The first concerns the problem of entry deterrence, especially the papers of Spence [26] and Dixit [4, 5]. These papers take the basic framework of Bain [1] and Sylos [27] and ask: What can the monopolist do prior to the entrant's decision point to make predation optimal in the short run? (The answers they give include such things as expanded capacity, sales networks, etc.) The relevance of this question is that the *threat* of predation is only credible if predation is *ex post* the optimal response, so the monopolist must make it so in order to forestall entry. What our model suggests (and what can be demonstrated formally) is that in repeated play situations, the actions taken by the monopolist need not make predation actually *ex post* optimal—what they must do is to make predation *possible* and, perhaps, increase the probability assessed by the entrants that it is *ex post* optimal. If deterrence is the objective, the appearance and not the reality of *ex post* optimal predation may be what is important.

The second context is that of a monopolist producer of a durable capital good where, for whatever reason, the monopolist is unable to maintain a rental market but must sell outright his product. In a multiperiod setting, where the monopolist is assumed to be sequentially optimizing, this can severely diminish the monopolist's market power. (See Bulow [2] and Stokey [25].) Supposing the monopolist produces subject to a capacity constraint, the monopolist is often *better off* with a tighter constraint. This is because the constraint prevents the monopolist from "over-producing." Then, if that constraint is the matter of private information for the monopolist, a monopolist with a loose constraint can successfully (in an equilibrium) masquerade as having a more stringent constraint, thereby recouping some of his lost market power. In essence, as the number of periods goes to infinity (as one comes closer to a continuous-time formulation), the monopolist can successfully attain the reputation of a "low capacity" producer even if his capacity is (with probability approaching one) high. Moorthy [15] presents an example along these lines.

ACKNOWLEDGMENTS

The authors express their gratitude to John Harsanyi, Mary Macgregor, Paul Milgrom, John Roberts, Robert Rosenthal, Reinhart Selten, Sylvain Sorin and many other colleagues for their helpful comments and criticisms. This research has been supported in part by National Science Foundation Grants SOC77-07741-A01, SOC75-21820-A01, SOC77-0600-A01 and SES80-06407 to the Institute for Mathematical Studies in the Social Sciences and the Graduate School of Business, Stanford University.

REFERENCES

1. J. S. BAIN, "Barriers to New Competition," Harvard Univ. Press, Cambridge, Mass., 1956.
2. J. BULOW, Durable goods monopolists, Stanford Graduate School of Business, mimeo, 1979; *J. Pol. Econ.*, in press.
3. V. CRAWFORD AND J. SOBEL, Strategic information transmission, University of California at San Diego, mimeo, 1981.
4. A. DIXIT, A model of duopoly suggesting a theory of entry barriers, *Bell J. Econ.* **10** (1979), 20-32.
5. A. DIXIT, The role of investment in entry-deterrence, *Econ. J.* **90** (1980), 95-106.
6. P. DYBVG AND C. SPATT, Does it pay to maintain a reputation, Financial Center Memorandum No. 32, Princeton University, 1980.
7. J. HARSANYI, Games with incomplete information played by Bayesian players, Parts I, II, and III, *Manag. Sci.* **14** (1967-1968), 159-182, 320-334, 486-502.
8. D. KREPS AND R. WILSON, "On the Chain-Store Paradox and Predation: Reputation for Toughness," Stanford University Graduate School of Business Research Paper No. 551, 1981.
9. D. KREPS AND R. WILSON, Sequential equilibria, *Econometrica* **50** (1982).
10. H. KUHN, Extensive games and the problem of information, in "Contributions to the Theory of Games, Vol. 2" (H. Kuhn and A. Tucker, Eds.), pp. 193-216, Princeton Univ. Press, Princeton, N.J., 1953.
11. M. MACGREGOR, A resolution of the chain-store paradox, University of California at Berkeley, mimeo, 1979.
12. P. MILGROM AND J. ROBERTS, "Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis," *Econometrica* **50** (1982), 443-460.
13. P. MILGROM AND J. ROBERTS, Predation, reputation, and entry deterrence, *J. Econ. Theory* **27** (1982), 280-312.
14. P. MILGROM AND R. WEBER, Distributional strategies for games with incomplete information, Northwestern University Graduate School of Management, mimeo, 1980.
15. S. MOORTHY, The Pablo Picasso problem, Stanford University Graduate School of Business, mimeo, 1980.
16. J. RILEY, Strong evolutionary equilibria and the war of attrition, *J. Theoret. Biol.* **82** (1980).
17. R. W. ROSENTHAL, Games of perfect information, predatory pricing and the chain-store paradox, *J. Econ. Theory* **25** (1981), 92-100.
18. R. W. ROSENTHAL, Sequences of games with varying opponents, *Econometrica* **47** (1979), 1353-1366.
19. R. W. ROSENTHAL AND H. J. LANDAU, A game-theoretic analysis of bargaining with reputations, Bell Telephone Laboratories, mimeo, 1979.

20. A. RUBINSTEIN, Strong perfect equilibrium in supergames, *Internat. J. Game Theory* **9** (1979), 1–12.
21. F. SCHERER, "Industrial Market Structure and Economic Performance," 2nd ed., Rand McNally College Publishing Company, Chicago, 1980.
22. R. SELTEN, Spieltheoretische behandlung eines oligopolmodells mit nachfragetragheit, *Z. Staatswissenschaft* **121** (1965).
23. R. SELTEN, Reexamination of the perfectness concept for equilibrium points in extensive games, *Internat. J. Game Theory* **4** (1975), 25–55.
24. R. SELTEN, The chain-store paradox, *Theory and Decision* **9** (1978), 127–159.
25. N. STOKEY, Self-fulfilling expectations, rational expectations, and durable goods pricing, Northwestern University, mimeo, 1979.
26. A. M. SPENCE, Entry, capacity, investment and oligopolistic pricing, *Bell J. Econ.* **8** (1977), 534–544.
27. P. SYLOS-LABINI, "Oligopoly and Technical Progress," trans. E. Henderson, Harvard Univ. Press, Cambridge, Mass., 1962.