RAND Journal of Economics
Vol. 31, No. 2, Summer 2000
pp. 207–236

A dynamic oligopoly with collusion and price wars

Chaim Fershtman*

and

Ariel Pakes**

We provide a collusive framework with heterogeneity among firms, investment, entry, and exit. It is a symmetric-information model in which it is hard to sustain collusion when there is an active firm that is likely to exit in the near future. Numerical analysis is used to compare a collusive to a noncollusive environment. Only the collusive industry generates price wars. Also, the collusive industry offers both more and higher-quality products to consumers, albeit often at a higher price. The positive effect of collusion on variety and quality more than compensates consumers for the negative effect of collusive prices, so that consumer surplus is larger with collusion.

1. Introduction

■ Most of the theoretical work on collusive behavior in oligopolistic markets assumes identical firms and/or an unchanging environment.[1] Useful as these assumptions are in clarifying both the process by which collusion can be supported and how it can break down, the framework needs to be modified before the alternative pricing schemes it generates will be used extensively by applied researchers.

Empirical researchers have constantly emphasized the extent of heterogeneity among firms within markets, and applied work is loath to assume that the different firms in a market have the same policy-cum-profitability options. This is particularly unfortunate because there is a very limited range of pricing models currently available

* Tel Aviv University and CentER for Economic Research; fersht@post.tau.ac.il.
** Harvard University and NBER; ariel@ariel.fas.harvard.edu.

We would like to thank two referees and an Editor for a set of detailed and useful comments, and the participants at seminars at Tel Aviv, Yale, Stanford, Harvard, Chicago, LSE, UCL, Cyprus, Lisbon, and Northwestern Universities for their many helpful suggestions. Paul McGuire provided excellent programing assistance. We are also thankful to the U.S.-Israel Binational Science Foundation and the NSF (grant no. SBR 9512106 through the NBER) for financial support. Ariel Pakes was at Yale University and visiting the University of Chicago while most of this research was done, and we thank those institutions for their hospitality.

[1] Collusion among asymmetric firms has been discussed in a number of articles (see, for example, Schmalensee (1987), Harrington (1989), Fudenberg, Levine, and Maskin (1994), and Compte, Jenny, and Rey (1999)). These articles differ from ours in that they assume the market structure is given and fixed over time.

to the applied researcher, and what there is does not conform terribly well to the movements in price vectors over time observed in a number of datasets.

Equally important is the fact that the assumptions of identical firms and an unchanging environment enable only a limited investigation of the implications of collusion; they only allow us to investigate the impact of collusion on prices. Whether an industry can or cannot support collusion also affects the incentives to launch products in (or to enter) the industry and to develop the products after they are launched. That is, the ability to collude will have an impact on the variety, cost, and quality of the products marketed by the industry, and this can have as much or more of an effect on welfare as do the price effects of collusion.

Much of the literature on collusion developed with an explanation for the price war phenomenon in mind. It was Stigler (1964) who first pointed out that price wars may be the outcomes of cheating on a collusive agreement, or of new entry into markets in which firms behave collusively. However, the first formalizations of these ideas in the literature on repeated games indicated that though the threat of reverting to price wars does play an important role in sustaining collusion, price wars did not exist on the equilibrium path (see Friedman, 1971; Rubinstein, 1979; and Abreu, 1986).

These early articles did make it clear that to analyze collusion we need to focus on the balance between the short-run gains from undercutting one's competitors (or deviating) and the expected long-run losses from the possibility of a breakdown in the collusive agreement caused by the deviation. By allowing for uncertainty and asymmetric information, Green and Porter (1984) were the first to obtain price wars as a part of equilibrium behavior. They considered repeated oligopolistic interaction with imperfect monitoring; the demand function was subject to random unobservable shocks, and firms did not observe their rivals' outputs. When a low price was observed, firms did not know if it was a consequence of a deviation from collusive pricing by one of their competitors or if there was a low realization of the demand shock. Green and Porter showed that some degree of collusion can be sustained in such games by trigger strategies that involve switching to price wars (punishment mode) whenever the price becomes lower than some endogenously determined threshold level. This result was later extended by Abreu, Pearce, and Stacchetti (1986), who considered the optimal cartel agreement in a repeated game with a general strategy space. The price war in the imperfect-monitoring model does not indicate a failure of the collusive agreement but is part of the equilibrium strategies designed to support the collusive outcome (see also Porter, 1983a, 1983b, and Ellison, 1994).

A second model with equilibrium price wars was introduced by Rotemberg and Saloner (1986); they modelled oligopolistic collusion with time-varying demand. In their setup, a period with high demand, a "boom," generates a greater temptation to deviate from the collusive agreement and hence countercyclical pricing.[2]

We follow these models in allowing firms to condition their quantity (or price) choices on the history of their interaction (thus allowing for collusive possibilities). We also allow, however, for investments to affect dynamic interactions. That is, in our model investment entry and exit processes allow firms to partially control the evolution of the vector of states that determine profits and consumer surplus. The incentives that underlie the investment decisions are determined by the nature of the collusive possibilities, while the ability to sustain collusion will depend on the vector of states that the investments have led to.[3]

[2] For further analysis of this model, see Bagwell and Staiger (1997), Staiger and Wolak (1992), Kandori (1991), and Haltiwanger and Harrington (1991). For a comparison of the two models and an empirical assessment of their applicability to the 1880 railroad cartel, see Ellison (1994).

[3] Davidson and Deneckere (1990) do provide a two-stage model of investment and collusion. However,

Since the outcome of the investment and entry processes is stochastic, over time the different firms will find themselves in different states facing different incentives. This will allow us to explicitly analyze the interaction between market structure and the ability to support collusion. It will also allow us to generate price wars (price vectors whose components all fall in response to a small change in structure) and to analyze the effects of collusion on welfare, taking account of the fact that collusive possibilities not only change prices conditional on achieved states but also change the distribution of states (the market structures) that are likely to be achieved.

We begin by modifying the framework presented in Ericson and Pakes (1995) to allow for collusion. This is a sequential model of oligopolistic interactions among a group of incumbents and potential entrants investing to explore profit opportunities. Firms' current states, together with the state of competitors from outside the industry, determine the current profits that result from any given price vector. Each period the incumbents engage in a pricing game that determines those profits and an investment game that determines the likelihood of tuples of future states. We follow Fershtman and Muller (1986) and consider a semicollusive industry in which firms may collude with respect to prices but play noncooperatively with respect to investment. Implicitly we are assuming prices are easy to observe while investment is not, and that this, together with the long-run and noisy nature of the outcomes of the investment process, makes collusion on investment decisions too difficult to support.[4] As in Ericson and Pakes, equilibrium is Markov perfect, but now strategies are allowed to depend on past pricing behavior, as well as on the "payoff relevant" states used in Maskin and Tirole (1988) and later in Ericson and Pakes.[5]

The model allows for exit from and entry into the industry. Exit rewards firms with a fixed selloff value for its equipment and occurs when the firm's continuation value is below this selloff value. The specification of the model ensures that in equilibrium each firm exits the industry in a finite number of periods with probability one. When a firm is near an exit state, we expect the market to have difficulty supporting collusion for two reasons. First, the ability to punish a firm that is near an exit state is limited. Second, the fact that one or more firms might exit provides an incentive for the continuing incumbent(s) to prefer noncollusive (lower-profit) current prices, since they will hasten the competitors' exit and leave the incumbent with fewer competitors, and hence higher profits, in the future.[6]

To allow for these phenomena in a realistic setting and still get fairly detailed results, we give up on the elegance of analytic results and rely instead on numerical analysis. In particular, we modify the computational algorithm developed by Pakes and McGuire (1994) to allow for collusion. To isolate the aspects of our results that are due to collusion, we compare the numerical results from an institutional structure that allows for collusive behavior to an industry with the same cost, demand, and investment

---

they assume all the investment takes place in the first stage and that there is no subsequent entry or exit, so that the second stage is a repeated game in which firms may differ in their capacity.

[4] For further analysis of semicollusive markets, see Davidson and Deneckere (1990), Fershtman and Gandal (1994), and Friedman and Thisse (1993).

[5] Note, however, that information remains symmetric, as in Maskin and Tirole (1997); for a discussion of extensions, see our Section 5.

[6] Of course, the exit of one competitor might be followed by the entry of another, but the fact that there are sunk costs and time required to enter ensures that the continuing incumbent will be temporarily better off. We note that it is assumed that at each market structure firms choose between a small number of easy-to-calculate price vectors (a collusive price, a noncollusive price, and a deviating price), and do not consider whether prices outside this set could be supported (see the discussion below).

primitives but with no collusive possibilities (perhaps because of an active antitrust authority).

Section 1 of the Sherman Act prohibits all contracts, combinations, and conspiracies in restraint of trade. Over the years the courts have used this act to take a clear position on collusive behavior such as price fixing. This behavior is per se illegal (it is considered illegal without having to prove that its effects are harmful). This position reflects the results of standard *static* economic analysis. Clearly, if we prohibit price fixing in a given period, consumer surplus in that period will rise (or at least not fall). However, the possibility of collusion and the type of collusive equilibrium affect firms' investment, entry, and exit decisions, as well as their pricing decision. That is, the market structure is determined, in part, by the collusive possibilities.

When we compare the equilibrium our model generates when we allow for collusion to the equilibrium that prohibits collusion, we find that the equilibrium with collusion generates a less concentrated market structure and offers both more and higher-quality products to consumers, albeit often at a higher price. Moreover, the positive effect of collusion on the variety and quality of products marketed more than compensates consumers for the negative effect of collusion on prices, so that consumer surplus is larger in the collusive environment. That is, our analysis shows that the presumption that collusion is necessarily bad for consumers is wrong, and this leads us to question the per se nature of antitrust policy toward collusion.

2. The model

■ We adapt the framework presented in Ericson and Pakes (1995) and the algorithm for computing it presented in Pakes and McGuire (1994) to allow for collusion. In each period there are n_t incumbent firms that differ in their physical characteristics, in say $\omega_{j,t}$. $\omega_{j,t}$ is the state of firm j at period t and evolves over time with the outcomes of an investment process. Positive outcomes in the investment process lead to states in which the firm makes higher profits. All investment decisions, including entry and exit decisions, are choice variables. Thus both the number of firms active, and their states, evolve as a controlled Markov process.

Decision making proceeds as follows. At the beginning of the period the incumbents decide whether to exit and potential entrants decide whether to enter. Entrants who do enter pay a sunk cost of entry and enter at a particular state in the *following* period (it takes one period to set up their plant and equipment). The incumbents who continue engage in a pricing and investment game. The pricing game sets each incumbent's price, say $p_{j,t}$, as a function of history. These prices, together with the firms' state variables, determine the profits of each active firm, say $\pi_{j,t}(\omega_t, p_t) \equiv \pi(\omega_{j,t}, p_{j,t}, \omega_{-j,t}, p_{-j,t})$, where $\omega_{-j,t} \equiv (\omega_{1,t}, \ldots, \omega_{j-1,t}, \omega_{j+1,t}, \ldots, \omega_{n_t,t})$ and $p_{-j,t} \equiv (p_{1,t}, \ldots, p_{j-1,t}, p_{j+1,t}, \ldots, p_{n_t,t})$.[7] Investments are directed at improving the firm's "physical" state, its ω value. The period concludes with the realizations of the stochastic outcomes of the investment and entry decisions.[8]

□ **Physical states, investment, and entry and exit.** As in Ericson and Pakes (1995), we assume that $\omega_{j,t}$ takes on values in the positive integers, $\omega_{j,t} \in \Omega \subset \mathcal{Z}^+$, so

[7] Note that our notation differs from that used in Pakes and McGuire. We index the state of the system facing the firm's decision maker by simply listing the firm's own state and the states of the firm's competitors.

[8] Though a change in the order of the moves would not change our characterization of behavior, it might change the actual numerical results. On the other hand, it would make no difference at all to what follows if we assumed that $\pi(\cdot)$ were expected profits, i.e., that realized profits could differ from it by a disturbance whose conditional mean is zero.

$\omega_t = (\omega_{1,t}, \ldots, \omega_{n_t,t}) \in \Omega^{n_t} \subset (Z^+)^{n_t}$. $\omega_{j,t}$ evolves over time with the outcomes of the firm's investment process, say $\eta_{j,t}$, and an industry-specific exogenous process that affects the profit opportunities facing the industry in a given period, say v_t. v_t captures the effects of factor prices and improvements in competition from products outside the industry—factors that generate positive correlation among the profits of our competing firms. Thus

$$\omega_{j,t+1} = \omega_{j,t} + \eta_{j,t+1} - v_{t+1}. \tag{1}$$

Both η and v will be nonnegative random variables, and the distribution of $\eta_{j,t+1}$ will be better, in the stochastic dominance sense, the larger is investment, our $x_{j,t}$. That is, the distribution of η is determined by the family

$$\mathcal{P} = \{P_\eta(\cdot|x), x \in \mathcal{R}_+\},$$

which is assumed stochastically increasing in x. The distribution of v is given exogenously.

If an incumbent decides to exit, it gets a selloff value of ϕ dollars and never reappears. We let $\chi_{j,t} \in \{0, 1\}$ indicate whether a firm exits ($\chi_{j,t} = 0$) or continues ($\chi_{j,t} = 1$).

Potential entrants decide whether to enter. To enter they must pay a sunk cost of x^e. An entrant appears in the following period as an incumbent at an $\omega = \omega^e \in \Omega^e \subset \Omega$ with probability p^e. For simplicity we assume that there is one potential entrant in every period, and we indicate whether entry occurs by the indicator function $\chi^e = \{0, 1\}$, $\chi^e = 1$ indicating entry.[9]

□ **Profits conditional on prices.** The version of the Pakes and McGuire (1994) algorithm currently available computes equilibria for different types of markets as determined by the profit function used in the calculations,[10] but Pakes and McGuire themselves provide a detailed numerical analysis of a dynamic differentiated product model. For explicit numerical results we will also need to work with a particular specification for cost and demand; to make it easy to compare our results to the noncollusive results available, we use the demand and cost structure used in Pakes and McGuire (however, as will become clear, we could have introduced collusive possibilities into any of the institutional structures computed by their algorithm).

Specifically, there are M consumers, each of which either chooses one of the $j = 1, \ldots, n_t$ goods in the market being studied or chooses to spend all of its income on the "outside alternative" (good 0). Consumer i who chooses good j obtains utility $U_{ij} = g(\omega_j) + (y_i - p_j) + \epsilon_{ij}$, where ω_j is an index of the quality of the product, $g(\omega_j)$ is the mean utility of consumers choosing good j (the average over consumers of the ϵ is zero for each j), p_j is its price, and y_i is the consumer's income. Since $g(0) = 0$, if the consumer chooses the outside alternative, its utility is $U_{i0} = y_i + \epsilon_{i0}$. Each consumer makes the choice that maximizes its utility.

[9] Note that the subset of ω at which the entrant enters is independent of the general progress of the industry, i.e., of the realizations of η and v. Thus entrants improve with the improvement of knowledge in the industry. If this did not occur, entry would eventually go to zero and stay there.

[10] A program to implement the algorithm can be accessed by FTP and is described in Pakes, Gowrisankaran, and McGuire (1995).

The function $g(\cdot)$ is introduced here simply to let us bound mean utility without resorting to a more complicated indirect utility function. Thus we set $g(\omega) = \omega$ if $\omega < \omega^*$, while if $\omega > \omega^*$ we choose $g(\cdot)$ to be increasing and have $\lim_{\omega \to \infty} g(\omega) \leq K$.[11]

To obtain the traditional logit form we assume that the $\{\epsilon_{ij}\}$ have independent (over both i and j) and identical type-1 extreme value distributions. Then the expected fraction of consumers who choose good j (where the expectation is taken over the ϵ), say $\sigma_j(\omega, p) \equiv \sigma(\omega_j, \omega_{-j}, p_j, p_{-j})$, is given by

$$\sigma(\omega_j, \omega_{-j}, p_j, p_{-j}) = \frac{\exp[g(\omega_j) - p_j]}{1 + \sum_{q=1}^{J} \exp[g(\omega_q) - p_q]}. \tag{2}$$

So if marginal cost is the constant c, then conditional on any set of prices, the profits of firm j are

$$\pi(\omega_j, \omega_{-j}, p_j, p_{-j}) = M\sigma(\omega_j, \omega_{-j}, p_j, p_{-j})[p_j - c]. \tag{3}$$

□ **Equilibrium.** A subgame-perfect equilibrium for the above game consists of a collection of strategies that constitute a Nash equilibrium for every history of the game. The strategies include price, investment and exit strategies for all incumbents, and entry strategies for potential entrants. We do not consider all such equilibria, only Markov equilibria that allow for collusive pricing arrangements enforced by punishment schemes. All strategies are allowed to depend upon both the "payoff relevant" physical states, the ω_t used in Maskin and Tirole (1988 and 1997), and on a set of indicator functions that keep track of whether any of the existing firms have ever deviated from a collusive pricing agreement in the past.

Formally we define the vector $\alpha_t = (\alpha_{1,t}, \ldots, \alpha_{n_t,t})$, with each $\alpha_j \in \{0, 1\}$, to indicate which, if any, of the existing firms have deviated in the past from the collusive strategies (in which case $\alpha_j = 1$). So the state of the system in period t will be characterized by the couple $(\omega_t, \alpha_t) \in \Omega^{n_t} \times \{0, 1\}^{n_t}$. The indicator function $\alpha_{j,t}$ evolves in a simple way: when a new firm enters the industry its $\alpha_j = 0$, and it becomes one only if the firm deviates. If the firm does deviate, its "α" stays at one for the remainder of its life.[12]

All strategies are assumed to be a function of the current value of the state vector, (ω_t, α_t). Consequently, a Markov-perfect equilibrium to our game is a tuple of strategies for the incumbent firms, $\{p_j(\omega, \alpha), x_j(\omega, \alpha), \chi_j(\omega, \alpha)\}$, and an entry strategy for the potential entrant, $\{\chi^e(\omega, \alpha) \in \{0, 1\}\}$, that constitute a Nash equilibrium at every (ω, α).[13]

Since we have already outlined investment, exit, and entry possibilities, to complete the specification of our model we need only a description of the pricing options facing a firm.

[11] More precisely, for $\omega > \omega^*$, $g(\omega) = \omega^*[2 - \exp - (\omega - \omega^*)]$, so the limit of $g(\cdot)$ is $2\omega^*$.

[12] One could modify this setup in several ways. For example, we could assume the indicator function stays at one for only a finite number of periods. This would, however, enlarge the state space and hence increase the computational burden of the problem. Alternatively, we could consider a stochastic punishment; in every period in the punishment phase there could be a fixed exogenous probability of ending the punishment and jumping back to the collusive mode.

[13] Since we use numerical methods to compute our equilibrium, what we are actually investigating is an ϵ Markov-perfect equilibrium (the computational procedure only ensures that the players are "nearly" optimizing). We, however, may make this ϵ as small as we wish.

□ **Pricing and collusive behavior.** Since we are allowing for asymmetric firms, even simple pricing rules generate prices that can be quite difficult to compute. Add to this the possibility of switching among different collusive pricing arrangements, and it is easy to see how the problem of determining equilibrium prices and the punishment scheme that can support them, and then coordinating which arrangement is to be used in every period, might be beyond the powers of any group of firms (especially firms that are worried about leaving a "smoking gun" the regulatory authorities might trace to them).

We thus limit our attention to an equilibrium with simple pricing possibilities. Specifically, we define a collusive, a deviant, and a noncollusive price, and we require that each of these prices can be computed from knowledge of the form of the current profit function (equation 3). We do, however, ensure that collusive prices will be used only when they can be supported by a punishment scheme. Since the deterrent value of the punishment must account for entry and exit, as well as for the fact that the states of incumbents vary over time, there is still a nontrivial computational problem in determining which states can support collusive prices.

Our model has the feature that the prices the firm chooses do not affect the evolution of the physical state, ω. This implies that one equilibrium of the pricing game is the Nash equilibrium to the static (one-shot) pricing game. The static Nash equilibrium ensures that each firm's price maximizes its current profits (given by (3)) given its competitor's prices and the characteristics of all products. As a result, these prices satisfy the vector of first-order conditions

$$-(p_j - c)\sigma_j(\omega, p)(1 - \sigma_j(\omega, p)) + \sigma_j(\omega, p) = 0; \qquad j = 1, \ldots, n_t, \tag{4}$$

where $\sigma_j(\omega, p)$ is given by (2). The unique solution to these first-order conditions (see Caplin and Nalebuff, 1991) will be denoted by $p^N(\omega)$, and the profits from this pricing rule and the current ω-tuple will be denoted by $\pi^N(\omega_j, \omega_{-j})$ (they are obtained by substituting p^N for p in (3)).

The equilibrium obtained when the pricing strategies $p^N(\omega_j, \omega_{-j})$ are always followed is the equilibrium that Pakes and McGuire (1994) analyze. We will compare it to an equilibrium that allows for collusive pricing but reverts to the one-shot Nash prices when collusion cannot be sustained. For this we need two other pricing rules: a collusive price to be denoted by $p^C(\omega_j, \omega_{-j})$ and the price a firm would charge if it were to deviate from the collusive pricing.

When all firms are identical, it is natural to focus on collusive arrangements in which the gains from collusion are distributed identically among firms, and the price maximizes total profits. There is less agreement on collusive rules when firms differ from one another and side payments are not possible, except perhaps for the conditions that the collusive agreement should increase all firms' profits and leave "better" firms better off (or at least firms with a higher "threat" value better off; for a discussion, see Schmalensee (1987)). A relatively simple solution to this problem that abides by these conditions is to assume that collusive prices are obtained as the solution to a (possibly implicit) bargaining game, and then use the Nash (1950) solution to that game.[14] Since we assume that when the collusive prices cannot be supported prices revert to $p^N(\omega)$ and profits to $\pi^N(\omega)$, we take the threat point for the bargaining game

[14] Related examples are often found in descriptive work. Thus when Scherer (1980) describes the cartel agreement in Germany during the 1920s and 1930s, he observes that since production capacity affected the firms' bargaining power, it determined market shares.

to be the one-shot Nash equilibrium payoffs, $\pi^N(\omega)$.[15] Thus the collusive pricing vector, say $p^C(\omega)$, is the solution to

$$\max_{p_1,\dots,p_{n_t}} \prod_{j=1}^{n_t} [\pi(\omega_j, \omega_{-j}, p_j, p_{-j}) - \pi^N(\omega_j, \omega_{-j})]_+, \tag{5}$$

where the + notation provides the positive part of the function (i.e., $[g]_+ \equiv \max[g, 0]$). Collusive profits as a function of the current ω-tuple are obtained by substituting p^C into the profit equation in (3) and will be denoted by $\pi^C(\omega_j, \omega_{-j})$.

The price an individual firm sets if it were to deviate from the collusive arrangement is the price that would maximize the deviant firm's profits given that the rest of the firms maintain collusive prices. That is, if the deviant firm's price is $p_j^D(\cdot)$, then it is obtained as the (unique) solution to

$$\max_{[p^D]} \pi(\omega_j, \omega_{-j}, p^D(\omega), p_{-j}^C(\omega)), \tag{6}$$

with deviant profits given by $\pi(\omega_j, \omega_{-j}, p_j^D, p_{-j}^C)$ (with $\pi(\cdot)$ from (3)).

We thus study an equilibrium with a very familiar pattern. Firms will choose the cooperative price $p^C(\omega)$ as long as none of the incumbent firms have deviated in the past and the collusive prices can be supported (see below). If there is ever a deviation by one of the incumbents, the firms revert to playing static Nash equilibrium prices, $p^N(\omega)$. Thus the strategies are in effect grim trigger strategies adapted to the dynamic setup that we study in this article. Since the environment changes over time, in every period in which collusion is possible (i.e., $\alpha_t = 0$) all incumbents must check to see if there are incentives for any one of them to deviate from the collusive agreement and play $p^D(\omega)$.[16] If there is a firm with an incentive to deviate, the collusive arrangement is not implemented and the firms revert to playing $p^N(\omega)$.

This completes the outline of our model. We have purposely kept it simple, perhaps too simple to be an adequate approximation to the institutional structure of any given industry. However, as we note, there are many modifications that can be made to our setup. Indeed, perhaps the biggest benefit of computational frameworks are that they can be perturbed to mimic many different institutions quite easily, hopefully enabling the researcher to map knowledge of the industry's institutions into the industry's collusive possibilities.

Our model does, however, have several features that are more realistic than most models used in the past to analyze collusion. In particular, we allow for heterogeneous firms whose states evolve over time according to the outcomes of an investment process, and we allow for entry and exit. We also allow for two types of strategic controls: prices and investments (defined to include entry and exit costs). Our base case assumes that firms can collude on prices provided the collusive prices can be sustained, but it does not allow for collusive investment rules. This reflects our belief that in many industries it is more difficult to discern and punish deviations from collusive investment

[15] We choose this "myopic" threat point, i.e., the per-period Nash equilibrium profits, for simplicity. An alternative that would make the analysis considerably more complicated would be to relate the threat point to the dynamic value of the Nash reversion. Actually, any procedure for selecting individually rational collusive prices in an environment that allows for asymmetric firms could be used here, and the appropriate procedure might well vary across industries.

[16] Note also that since, conditional on the state, *any* deviation is followed by the same punishment, the deviating price $p_j^D(\cdot)$ is indeed the optimal deviation strategy.

policies than from collusive pricing policies. We will, however, present numerical results that compare this base case to two other possibilities, one in which there are no collusive possibilities for either prices or investment, and one in which there can be "perfect" collusion on both strategic variables (by perfect we mean that we need not worry about enforcement constraints, as would be the case if the industry were run as a multiproduct monopoly).

3. Computing the equilibrium

■ This section modifies the iterative computational algorithm provided in Pakes and McGuire (1994) to allow for our model of collusion.

As in that article, we assume that ω resides on the integers and invoke arguments analogous to those in Ericson and Pakes (1995) to show that in equilibrium,

(i) We only observe ω values in a finite set, say $[1, \ldots, \overline{\omega}]$, and
(ii) We never observe more than a finite set of firms active, say $\mathcal{N}$.

It follows that if we let $\omega_j = 0$ indicate that the jth firm is not active, then the observed ω-tuples take values in $\Omega^{\mathcal{N}}$, where $\Omega = [0, 1, \ldots, \overline{\omega}]$.

Recall that in our model behavior depends both on the ω-tuple *and* on whether any of the currently active firms have ever deviated from a collusive agreement in the past, on $\alpha \in \{0, 1\}^{\mathcal{N}}$ $[\alpha_j \equiv 0$ if $\omega_j = 0]$. Thus the "state" of our industry is fully described by the couple $(\omega, \alpha) \subset \Omega^{\mathcal{N}} \times \{0, 1\}^{\mathcal{N}}$.

We also use Pakes and McGuire's adaptation of the Ericson-Pakes model for the transitions of ω for the firms that remain active. Since ω_j can be interpreted as a difference between the quality, or average utility, of the good marketed by firm j and the mean value of the outside alternative, differences in the value of ω over time are a result of the difference between the outcomes of the firm's investments and any exogenous increments in the value of the outside alternative.[17] We assume that in each period the state of the quality both of the firm's product and of the outside alternative can move by at most one unit, and we let $p(x)$ be the probability that the firm's product improves and δ be the probability that the outside alternative improves. Then $\omega_{j,t+1} = \omega_{jt} + 1$ with probability $p(x)[1 - \delta]$, $\omega_{j,t+1} = \omega_{jt} - 1$ with probability $[1 - p(x)]\delta$, and $\omega_{j,t+1} = \omega_{jt}$ with the remaining probability.

Firms that invest more have a larger chance of improving their ω, so we require $\partial p(x)/\partial x \geq 0$. Moreover, we assume both that $p(0) = 0$ (so that a firm cannot improve its ω without investment) and that $p(\cdot)$ is concave in x (this makes it easier to solve for the optimal x). A specification for $p(x)$ that satisfies all these conditions is $p(x) = ax/(1 + ax)$, and we use this specification in the numerical calculations.

We compute the equilibrium strategies using the value function approach (for more details, see Starr and Ho (1969)). Any suggested strategies will be Markov-perfect equilibrium strategies if they are optimal given the value function, and the value function is the continuation value of the game if the firms indeed play those strategies. We then show how that value function, and the associated strategies, can be computed.

□ **The Bellman equation.** We compute an equilibrium in which no firm ever deviates, but the values generated by deviant behavior (i.e., by behavior "off the equilibrium path") determine when the industry can support collusion and hence must be

[17] Note that the realizations of υ cause positive correlation in the demand, and hence the profits, of the firms in the industry. Without the υ the model would predict a negative correlation among their profits, a prediction at odds with the data on the evolution of most industries (see Pakes and McGuire (1994) for more details).

computed. Since the punishments from deviating are determined by values in which only one firm has deviated, it suffices to compute values for $(\omega, \alpha) \in S$, where

$$S \equiv \{(\omega, \alpha) \mid \alpha \in \{0, 1\}^{\mathcal{N}} \quad \text{and} \quad \sum \alpha_j \in \{0, 1\}, \omega \in \Omega^{\mathcal{N}}\}.$$

We first present the Bellman equation, and then we explain it. For each $(\omega, \alpha) \in S$, the value of the jth firm satisfies

$$V(\omega_j; \omega_{-j}, \alpha) = \max\Biggl\{\phi, \pi(\omega_j; \omega_{-j}, \alpha) + \max_{x \geq 0}\Bigl[-x + \beta \sum_{\omega'} V(\omega_j'; \omega_{-j}', \alpha')p(\omega_j' \mid \omega_j, x)p(\omega_{-j}' \mid \omega, \alpha)\Bigr]\Biggr\}, \tag{7}$$

where

$$\pi(\omega_j; \omega_{-j}, \alpha) \equiv I(\omega, \alpha)\pi^C(\omega_j; \omega_{-j}) + [1 - I(\omega, \alpha)]\pi^N(\omega_j; \omega_{-j})$$

and

$I(\omega, \alpha) \in \{1, 0\}$, with $I(\omega, \alpha) = 1$ if and only if

(i) $\alpha = 0$, and
(ii) for all j who continue,

$$\pi_j^D + \max_{x_j}\Bigl[-x_j + \beta \sum_{\omega'} V(\omega_j'; \omega_{-j}', \alpha')p(\omega_j' \mid \omega_j, x_j)p(\omega_{-j}' \mid \omega, x, \alpha_j = 1, \alpha_{-j} = 0)\Bigr]$$
$$< \pi_j^C + \max_{x_j}\Bigl[-x_j + \beta \sum_{\omega'} V(\omega_j', \omega_{-j}', \alpha')p(\omega_j' \mid \omega_j, x_j)p(\omega_{-j}' \mid \omega, x, \alpha = 0)\Bigr]. \tag{8}$$

The first max operator compares the exit value of the firm (ϕ) to its continuation value. If ϕ is larger, the firm shuts down. We let $\chi(\omega_j; \omega_{-j}, \alpha)$ be the indicator function, which takes the value of one if the firm remains active and zero elsewhere.

If the firm does continue, it earns current profits plus the expected discounted value of future returns. Current profits are either Nash profits or collusive profits according to whether the indicator function, $I(\cdot)$, is zero or one. If $\alpha \neq 0$, that is, if one of the current participants deviated in the past, then $I = 0$ and the firms earn the one-shot Nash equilibrium profits. If $\alpha = 0$, then the firms collude only if collusion can be sustained; that is, if for every active firm collusive profits plus the expected discounted value of future net cash flows conditional on the firm colluding are greater than defector's profits plus the expected discounted value of future net cash flows conditional on the firm defecting.

Note that the distribution of the firm's future competitors, ω_{-j}', depends on entry and exit decisions made this period. We have already formalized exit. As noted, we allow for one potential entrant in every period. If the entrant enters, it must incur (sunk) setup costs of x^e ($> \phi$) dollars and spend a period building its plant. In the subsequent period it becomes an incumbent with $(\omega_j, \alpha_j) = (\omega^e, 0)$ with probability p^e

(the set of ω^e and their probabilities, p^e, are exogenously specified). The entrant enters if it is profitable to do so, that is, if

$$\beta \sum_{\omega'} V(\omega_e; \omega'_{-e}, \alpha')p^e(\omega^e)p(\omega'_{-e} | \omega, \alpha) - x^e > 0. \tag{9}$$

If this condition is satisfied we set $\chi^e(\omega, \alpha) = 1$, otherwise $\chi^e(\omega, \alpha) = 0$.

Note also that α' is determined by ω' and α. If $\alpha = 0$, then, since neither incumbents nor entrants deviate on the equilibrium path, $\alpha' = 0$. If $(\alpha_j, \alpha_{-j}) = (1, 0)$ for some j, then in states where the ith firm remains active, α' has a one in the ith slot and zero elsewhere, whereas if the ith firm exits, $\alpha' = 0$.

Finally we adopt the convention that in each period exit decisions are made first, followed by entry decisions, and then investment decisions. Firms that exit are assumed not to make profits in the current period. Since entrants do not make profits either, profits are calculated from the ω-tuple of the incumbents that remain active.

If there is a function $V(\cdot)$: $S \to R$, and a set of investment, exit, entry, and collusion rules that satisfy the Bellman equation pointwise for each $(\omega_j, \omega_{-j}, \alpha) \in S$, then those policies, by construction, are equilibrium policies to our game. That is, given these policies the value function $V(\cdot)$ provides the (expected discounted) value of all incumbents and potential entrants, and given this value function the policies are optimal at every history of the game.

□ **Computing the fixed point.** This subsection is designed to provide the reader with one algorithm for computing the collusive equilibria (the fixed point to the Bellman equation in (7)).[18]

The computational algorithm is iterative. Temporarily assume S, i.e., $\overline{\omega}$ and $\mathcal{N}$, are known. Given $(\overline{\omega}, \mathcal{N})$ we explain what is in memory at each iteration and then show how these objects are updated. We then come back to the problem of finding $(\overline{\omega}, \mathcal{N})$.

Letting z^k denote the kth iteration's estimate of z, going into iteration $k + 1$ we have in memory

(i) for each $(\omega, \alpha) \in S$

- (a) the prior iteration's estimates of the value function and the investment policies, say $(V^k(\cdot), x^k(\cdot))$, and
- (b) $\pi^N(\cdot)$,

(ii) $\pi^D(\cdot)$, $\pi^C(\cdot)$ for the subset of S in which $\alpha = 0$.

To go from iteration k to iteration $k + 1$ we need to update our estimates of $(V(\cdot), x(\cdot))$. To do so we cycle through the points in S in a predetermined order, performing the following sequence of calculations.[19] We start the calculations at the current point by determining its entry policy for iteration $k + 1$. This is done by substituting the appropriate components from $(V^k(\cdot), x^k(\cdot))$ held in memory for the V's and x's

[18] There are a number of ways one might modify the procedure we introduce here, several of which are discussed in more detail for the analogous problem in Pakes and McGuire (1994). That article also formally introduces the operators that define the algorithm (a step that seemed unnecessary here, given that it has already been done once).

[19] Note that the calculations at each point use the current iteration's estimates of $V(\cdot)$, $x(\cdot)$ for the points ordered before the point being currently updated; i.e., we use Gauss-Seidel iterations and not the method of successive approximations.

appearing in equation (9), and then setting $\chi_e^{k+1}(\cdot) = 1$ if that condition is satisfied and zero otherwise.

Next we determine the collusion rule for this iteration at the given point. This is done by using the information in memory and $\chi_e^{k+1}(\cdot)$ to determine whether $I^{k+1}(\omega, \alpha)$ in the Bellman equation in (7) should be set to one (in which case prices are the collusive prices) or zero. If $\alpha = 1$ (somebody has deviated in the past) the firms do not collude and we set $I^{k+1}(\cdot) = 0$. If $\alpha = 0$ we substitute the appropriate elements of $(V^k(\cdot), x^k(\cdot))$ from memory for the V's and the x's in (8) and then check whether they imply that we can support the collusive equilibrium (in which case $I^{k+1}(\omega, \alpha) = 1$).

Given the entry and collusion rules, the $k + 1$th iteration's investment policy for each incumbent at the given point is determined to

$$\max_{x \geq 0} \left[-x + \beta \sum_{\omega'} V^k(\omega_j'; \omega_{-j}', \alpha') p(\omega_j' \mid \omega_j, x) p(\omega_{-j}' \mid \omega_{-j}, x_{-j}^k, \chi_e^{k+1}) \right], \tag{10}$$

where $p(\omega_{-j}' \mid \omega_{-j}, x_{-j}^k, \chi_e^{k+1})$ is notation for the probability of ω_{-j}' given $\chi_e^{k+1}(\cdot)$ and that the incumbent's competitors invest the amount determined in the prior iteration, i.e., $x^k(\cdot)$. Note that $x_{-j}^k = x_{-j}^k(\omega, \alpha)$ and differs if an incumbent has deviated in the past.

This calculation produces both an x_j^{k+1}, which is the investment policy copied directly into memory, and an ability to calculate the continuation value for firm j, the value

$$\begin{aligned} V_j^{k+1}(\cdot) = {} & \pi_j[\cdot, I^{k+1}(\cdot)] - x_j^{k+1} \\ & + \beta \sum_{\omega'} V^k(\omega_j'; \omega_{-j}', \alpha') p(\omega_j' \mid \omega_j, x_j^{k+1}) p(\omega_{-j}' \mid \omega_{-j}, x_{-j}^k, \chi_e^{k+1}). \end{aligned} \tag{11}$$

If this value is greater then the selloff value of the firm (ϕ), we set χ_j^{k+1}, the indicator function for whether the firm remains, to one, otherwise this indicator function is set to zero.

If none of the incumbents exit, we copy the continuation values in (11) into memory, thus completing the sequence of calculations for the current point for this iteration. If one or more of the incumbents have continuation values less than ϕ, we set those firms' V_j^{k+1} to ϕ and copy over all policies (entry, collusion, investment, and exit policies) from the (ω, α) point obtained from the current (ω, α) by resetting the (ω_j, α_j) couple for the exiting firm(s) to zero.

That completes an iteration. We continue iterating until norms of the difference between the estimates of *both* the value function and the investment policies calculated at successive iterations are below some critical value.[20] It is straightforward to check that if $(V^{k+1}(\cdot), x^{k+1}(\cdot)) = (V^k(\cdot), x^k(\cdot)) \equiv (V^*(\cdot), x^*(\cdot))$, then $(V^*(\cdot), x^*(\cdot))$ and the associated entry, exit, and collusion policies (all of which are uniquely determined by $(V^*(\cdot), x^*(\cdot))$ satisfy the fixed-point condition in the Bellman equation.

[20] For a discussion of stopping rules, see Judd (1998). We used a Euclidean norm normalized by the standard deviations at a given iteration to measure the difference between successive iterations. The normalization is to ensure that differences between iterations are independent of the units of measurement used. Since we started with an overestimate of the value function, we also found it useful to keep track of differences between the means of successive iterations. The algorithm stopped when both norms were within 10^{-5} of zero *and* the difference between successive iterations' means were less than 10^{-5}.

Note that we use last-iteration's investment policies to calculate both the entry and the collusion rules and to form the expected discounted value of the future that determines the incumbents' new investment values. This simplifies the calculations at each point significantly. In particular, with these substitutions at each point we need only

(i) check the entry condition
(ii) check the collusion condition
(iii) solve a set of single agent optimization problems (for the investment of incumbents), and
(iv) check for exit.

Had we attempted to obtain a simultaneous solution for the entry, collusion, exit, and investment rules that are consistent with the current iteration's value function, we would have had to solve a simultaneous equations system at each point at each iteration. Of course at the fixed point the current iteration's policies are identical to those of the last iteration, so the solution for the policies *is* a simultaneous solution.

We now come back to the problem of obtaining S. As in Pakes and McGuire (1994), the 0 and $\overline{\omega}$ that determine Ω are taken from the solution to the monopolist's problem, i.e., by the point at which the monopolist exits and the point at which the monopolist stops investing, respectively.[21] $\mathcal{N}$ (the maximal number of firms ever active) is determined by first computing the equilibrium by limiting the number of firms ever active to some small number, and then gradually increasing that number until there are no equilibrium points for which there are $\mathcal{N}^*$ firms active and at which a potential entrant would want to enter. Then $\mathcal{N} = \mathcal{N}^*$.[22]

4. Numerical results

■ We begin with a description of the parameters of the problem together with a short description of the industry equilibrium when collusion is not allowed. We then describe the entry, exit, investment, and collusion policies in the collusive equilibrium. The last section provides the descriptive and normative statistics we obtained when we simulated our market's evolution twice, once allowing for collusion and once not.

□ **Parameter values and the noncollusive equilibrium.** The precise parameter values used to obtain numerical results are similar to those in Pakes and McGuire (1994) (an article that did not allow for collusion). We did decrease the market size parameter (M) by a factor of five (from $M = 5$ to $M = 1$) and increase the entry costs by a factor of 4 (from $x^e = .2$ to $x^e = .8$). The scrap value at exit (ϕ) stays at .1, so that the unrecoverable entry costs go up from .1 to .7. Entry costs are now about $\frac{1}{6}$ of total production costs within a period, compared to about $\frac{1}{125}$ in Pakes and McGuire. As in that study, the entrant who pays the sunk costs enters in the following period at $\omega = 4$,

[21] It can be shown that an ω low enough to induce the monopolist to exit will induce exit from any alternative market structure. Ericson and Pakes (1995) show that the boundedness of the value function implies that there exists an ω above which no firm invests. Since without investment a firm cannot improve its ω, this puts an upper bound on the observed ω's. Numerically, we have found that the ω at which the monopolist stops investing is higher than the ω that halts investing in any other market structure (though we have no proof that this must be so). The program checks that there is no equilibrium point at which a firm with an $\omega = \overline{\omega}$ actually wants to invest.

[22] We note that as in Pakes and McGuire there is no guarantee either that the algorithm will converge or that, given convergence, the equilibrium is unique. We have run into periodic convergence problems, and these are treated as in Pakes and McGuire. As noted, it is clear that the Nash equilibrium without collusion is an equilibrium, but if we started out with high-enough initial conditions, we always converged to the one equilibrium reported in the text.

if the outside alternative does not move up in the interim, and at $\omega = 3$, if it does move up.[23]

Pakes and McGuire's base case had the maximum number of firms ever active $\mathcal{N} = 6$, and it generated an industry in which it is relatively cheap to start up and explore a new idea. Most startups were unsuccessful (had negative realized discounted values), but the few that were successful tended to earn phenomenal rates of return on their investment. The changes we have made in the parameters make the industry smaller (now $\mathcal{N} = 4$ when we don't allow for collusion), and entry and exit are much less frequent. So we introduce the possibility of collusion in a fairly stable environment, one where one might think collusion could be reasonably successful.

For our parameter values the noncollusive industry is largely a "natural monopoly." Simulating 100,000 periods of industry evolution, we found that in 90% of the periods there was only one firm producing. The quality of the product of the monopolist did eventually fall low enough relative to advances in the outside alternative to induce entry. Entry was followed by a period of active investment competition among firms with low-quality products. The competitive phase was very unstable; there was a third entrant if the investments of the first two were not successful, and it was quite likely that one or both of the firms that were initially competing, and some of the subsequent entrants would fail (the modal life span is two years and the median is three). This competitive phase typically ended with one firm pulling away from the others, and the others finding it unprofitable to compete with the superior product and exiting the industry.

Entry only occurs in such an industry when the incumbent firms have a relatively low ω. For example, when there are two active firms in the industry, entry by a third firm occurs only when both active firms have $\omega \leq 4$. Incumbents invest more when their states are more similar, since this is when competition for developing a successful high-quality product is most intense.

□ **The collusive industry: policy output.** *Collusive states.* Figure 1 looks at three firm states; it first asks which of the three firms will continue and which will exit, and then it plots the borders of the region in which we will observe collusion among the continuing firms. The origin is the point hidden behind the intersection of the tunnels, and a firm typically exits when its ω is two.

The continuing firms *cannot* support collusion inside the three tunnels along the axis, but they can support collusion everywhere else. The inside of each of these tunnels represents a set of points in which two of the three firms are near an exit state. The diagram does not have a raised "floor" in each orthant. That would represent a region in which two of the three firms have relatively high ω's, and collusion breaks down because of a low ω of the third firm. This is because when two of the three firms have a high ω and the third firm does not exit (its $\omega \geq 3$), the three firms always collude.

Thus we only observe breakdowns of collusion among three continuing firms when *two of the three* firms have ω's that are near (but still above) the exit state. Just how low depends on how high the ω of the leading firm is. The higher the leading firm's ω, the smaller the area in which collusion can be supported. Thus collusion can be

[23] The other parameters are set at $\beta = .9$, $\delta = .65$ (δ is the probability that the outside alternative moves up), and $c = 5$. We note that these parameters produced an $\overline{\omega} = 25$, and $\mathcal{N} = 4$. Lemma 4 in Pakes (1994) and the definition of S in Section 3 then imply that S contains about 100,000 points. When there were no convergence problems, a run would take about four hours on our Sun Sparc 2 workstation.

FIGURE 1

BOUNDARIES OF COLLUSION REGION

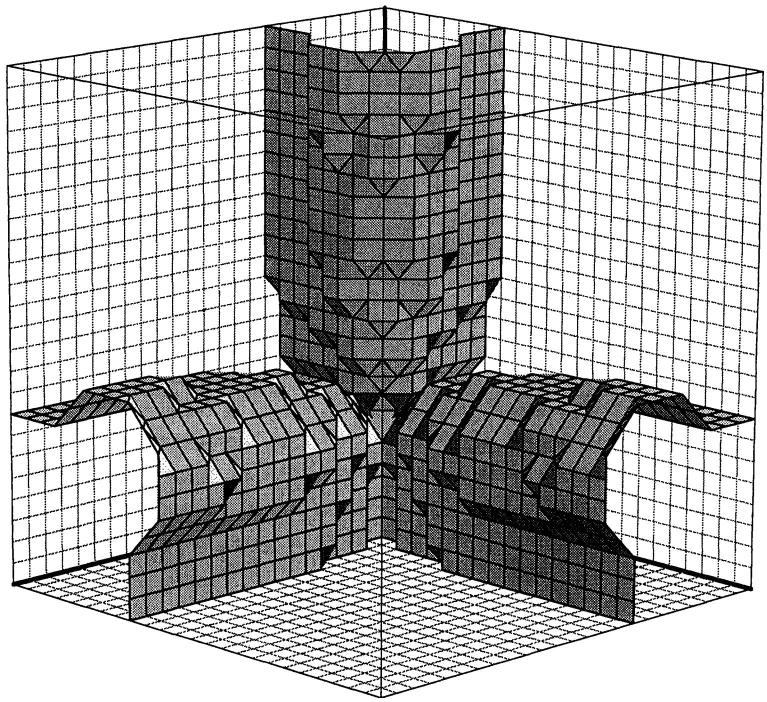

supported at (8, 6, 6), but it cannot be supported at any ($\omega \geq 9$, 6, 6), etc. As we shall see, the higher the ω of the leading firm, the more likely that firm is to outlast its weaker competitors and eventually obtain a dominant monopoly position.

The orthants of this figure provide the collusive regions when two firms are active. Collusion breaks down when one of the two is near an exit state, but again just how near depends on where the larger competitor is. When the two firms have the same value of ω, then they will collude as long as that $\omega \geq 5$. On the other hand, when their ω's differ, collusion can break down even when both values of ω are much higher than that. Thus collusion cannot be supported at ω couples equal to either (5, $\omega \geq 6$) or (8, $\omega \geq 12$) (below we see more clearly the collusive region when two firms are active). Though $\mathcal{N} = 4$, the results below indicate that we will rarely observe periods with $n = 4$, so we ignore those states here.

There are two reasons to think that it might be harder to support collusion when some of the continuing incumbents are close to states that would induce exit. Both have to do with the fact that all players realize that in these states the smaller firms have a good chance of exiting over the next few periods.

Insufficient punishment. The small firm(s) realizes that the other market participants may not be able to punish it severely enough if it deviates, and this provides it with an incentive to deviate.

"Predatory" behavior. The large firm realizes that the future will be better if it can "force" the small firm(s) to exit and then monopolize the market until the next entrant arrives. Since the smaller firm is more likely to exit if there is no collusion,

the larger firm has an incentive to deviate from the collusive pricing in order to increase the likelihood of the smaller firm exiting.[24]

To distinguish between these two reasons for not being able to support collusion, we looked at the policy output at each point in the set of two- and three-firm equilibria where collusion breaks down and checked which firms preferred to deviate. In most states where collusion broke down, all firms preferred to deviate. However, at states near the border of the collusive region, we found that frequently the smaller firm(s) would have preferred a collusive arrangement; only the large firm would have deviated (indeed, the largest firm preferred to deviate at every point at which collusion could not be sustained). On the other hand, as one might have guessed, in three-firm equilibria the middle-sized firm had the least incentive to deviate; i.e., if two firms preferred to deviate, it was always the largest and the smallest firm.

This is an industry that is extremely profitable to a monopolist with an even modestly large ω. The relatively large entry costs imply that a monopolist with a moderately large ω can deter entry, so it is likely that the monopolist can maintain its monopoly position for some time. As a result, when there are a number of active incumbents and one draws ahead of its competitors, that incumbent finds itself better off acting "predatorily," and increasing its probability of becoming a monopolist, than colluding.[25] On the other hand, whenever the states of two or more incumbents are sufficiently high, monopoly becomes too remote a possibility, and collusion can be supported.

In models that allow for exit, there is a "predatory" reason for not colluding, and we find that it is often the reason collusion cannot be supported. We investigate a particular model that allows for such behavior, but it is clear that many of our detailed assumptions could be changed and this difficulty in supporting collusion would persist.[26]

One final point about the collusive region illustrated in Figure 1: Recall that if two firms have high-enough states, they will collude even if there is a third firm at a much lower state ($\omega_3 = 3$ *or* 4). Since we have set the states at which a potential entrant will enter if it should desire to do so (our ω_e) at 3 or 4, this implies that there is likely to be collusion after entry provided the incumbents' states are high enough—a fact that helps produce a rather complex set of entry regions.

Entry states. Figure 2 provides a three-dimensional plot of the states at which entry occurs (the shaded area). There is an entry "hill" at the origin and a single entry "plateau" in each of the other orthants (as explained below, each of the plateaus actually has two levels). The plateaus correspond to entry states in which multiple firms

[24] When the firms are not colluding, the market prices are the Nash equilibrium prices, so none of the firms ever charges a price below marginal cost. Thus the use here of the phrase "predatory" is meant only to be indicative of the fact that the leading firm wishes to break the collusive agreement to induce its competitors to exit and not to predatory pricing, as the term is commonly used in the Industrial Organization literature.

[25] Relatedly, when only a single firm is active, if that firm were given the option of being labelled a "deviant" firm, it would take it. This is because an incumbent that is a deviant does better at deterring future entry. See also the discussion in the next footnote.

[26] For example, we could assume that punishment phases last only a finite number of periods, and/or terminate when there is only one firm in the market (which would eliminate the advantage to deviating that results from an ability of the single firm to deter future entry). Alternatively, we could allow the incumbent firms to revert to the Nash price equilibrium for a fixed number of periods immediately following entry, thus "fighting" instead of "accommodating" the entry (though this might risk a predatory pricing investigation). Though our techniques can be adapted to investigate such alternatives, we do not do so here.

FIGURE 2

ENTRY REGION

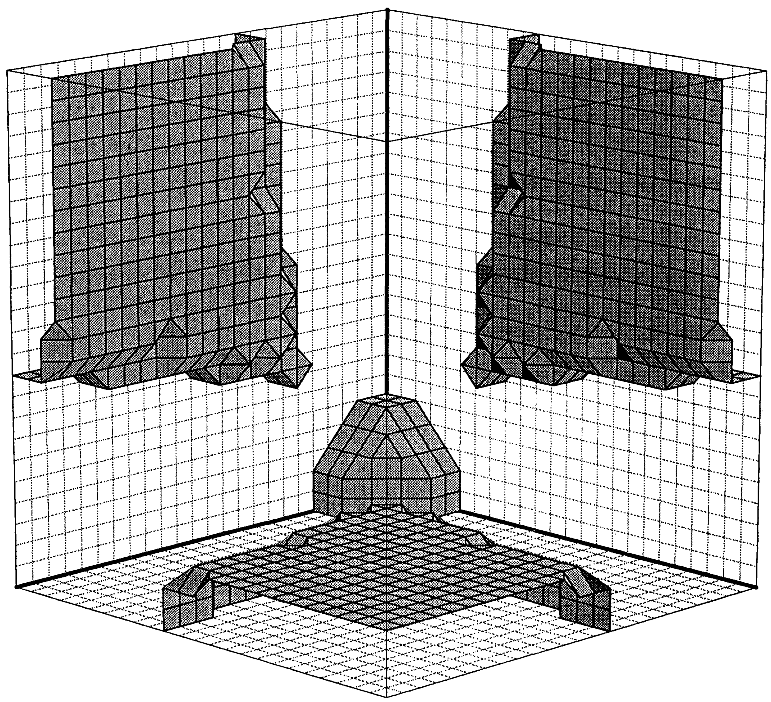

are active and at least two of their ω's are quite large. The hill at the origin corresponds to entry states in which all active firms have very low values of ω.

The discussion of collusion makes it clear why we should expect the plateaus. When there are multiple firms active and they all have high values of ω, then, as illustrated in Figure 1, an entrant knows that the incumbents will collude with it after it enters. The entry question is then a question of whether the inducement generated by the possibility of collusive profits will be sufficient to generate entry. It will be if there are only two firms active. But when there are three firms active, then the inducement generated by collusive profits will generate entry only when the third incumbent has an $\omega = 3$ (the third incumbent would exit were its ω to drop to two). Thus the upper "tier" of the entry plateau corresponds to entry with three firms active with one of those firms at an $\omega = 3$, and the lower tier corresponds to entry with only two (large) incumbents.

Entry also occurs when *all* incumbents have *relatively low* ω values. In this case, the entrant enters despite the fact that it will not earn collusive profits because it has a reasonable probability of becoming either a large dominant player in the future or one of a small set of profitable future colluders (see below).

Entry does *not* occur when two firms are active when those firms have only moderately high ω's; i.e., ω's larger than the likely postentry states of the entrant, but not high enough for collusion to be supported after entry. Note that this implies that entrants *prefer* an industry with very strong incumbents to one where incumbents are only modestly strong, as collusion can be sustained in the former but not the latter case. So entrants prefer industries with either very strong or very weak incumbents.

Finally, when only one firm is active it can deter entry provided its ω is at least moderately large (≥ 5). This is because the potential entrant knows that were the entrant to enter, the larger incumbent would both price "predatorily" and invest heavily. This plus the relatively large entry costs deters entry.

The contrast between the entry states in the model that allows for collusion and the entry states in the model that does not is quite striking. The plateaus simply *disappear* when collusion is not allowed. Further, without collusion, when two firms are active entry occurs only when both active firms have an $\omega \leq 4$. Therefore there are also more low states at which entry occurs when there is the possibility of collusion in the future.

As a result, we should expect to see more firms active when we allow for collusion than when we do not. Indeed, if in a situation that allows for collusion, the profits at any given tuple of states when collusion breaks down are at least as high as the profits that would be earned at the same tuple of states in an institutional setting that does not allow for collusion, we would always expect more entry when collusion is allowed. This is a fact that has largely been ignored in the literature and has important implications for the welfare analysis of collusive behavior. Interestingly, it is also consistent with empirical findings. Thus Symeonidis (1999) concludes that the introduction of laws to restrict cartels in the United Kingdom in the late 1950s "had no significant effect on profits, while it had a strong negative effect on the number of firms" (p. 1).

The value function and investment. Figures 3 and 4 are plots of sections of the value function of the firm. Figure 3 assumes two firms are active: the firm we are studying (which has an ω that increases as we move away from the reader) and its competitor (whose ω increases as we move to the right on the graph). Figure 4 assumes three firms are active: the two whose ω's are plotted on the axis and a third firm at $\omega = 3$. Figure 5 provides the probability that our firm's research is successful when only two firms are active (recall that this is a monotone transform of the firm's investment expenditures), while Figure 6 provides the same probability but this time when there is a third firm active and its $\omega = 4$. Since the role of investment in this model is to provide a larger probability of moving up the value function, the relationship between the two figures is that the slope of the value function determines the level of investment. The darkly shaded areas in these figures provide the states at which another firm would enter; the lightly shaded area, when added to the entry area in the northeast portion of the figure, provides the states at which collusion can be sustained.

These figures reinforce some of our earlier remarks. Thus they show that collusion is more likely to be sustained when the firms are more similar, and that entry will occur when the incumbents have either very low or very high values of ω, but not when their ω's are in an intermediate range. They also show that there can be entry when there is a third firm and its $\omega = 3$, but if the third firm has an $\omega = 4$ or more, and the other incumbents have moderately high states, there will be no further entry.

A few other points also come out clearly. First, the firm's value, though monotone in its own ω, is not monotone in its competitor's state. The firm clearly prefers to be the only firm with a high ω, for then it can force its competitor to exit and charge monopoly prices until another entrant appears. But if its competitor is moderately large, the firm would prefer that the competitor become larger yet, as this would enable them to sustain collusion. That is, the firm's value function actually increases in the value of its competitor's state in the region bordering the collusive states. Once both firms are in the collusive region, the value function looks relatively flat (though, as we shall presently see, there are small "hills" in this region that don't become apparent with the resolution in this figure). When there is collusion, the benefits from a firm increasing

FIGURE 3

VALUE FUNCTION, $\omega_3 = 0$

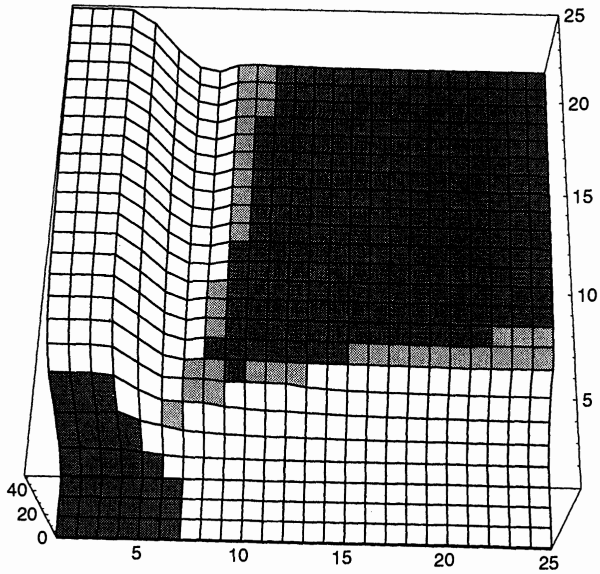

the quality of its product are shared among the incumbents, and sometimes with a new entrant also. The result is a value function that flattens out rather quickly after entering the collusive region.

There is one clear difference between Figures 3 and 4. In Figure 3 (which, recall, assumes only two firms active), there is only a narrow band of states in which the two

FIGURE 4

VALUE FUNCTION, $\omega_3 = 3$

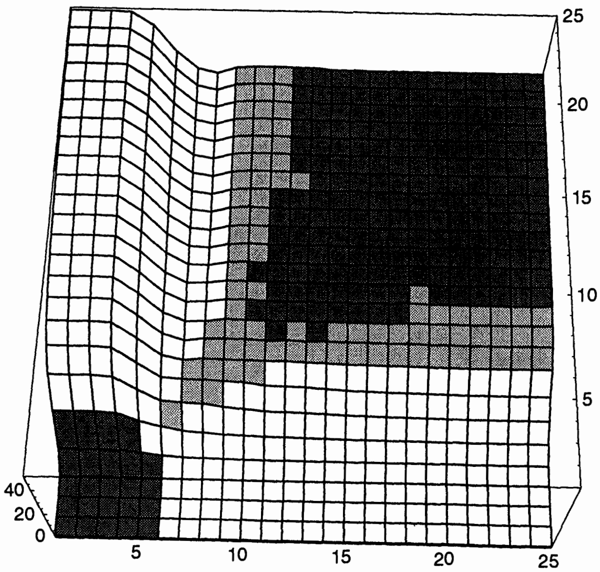

FIGURE 5

INVESTMENT (PROBABILITY OF SUCCESSFUL RESEARCH), $\omega_3 = 0$

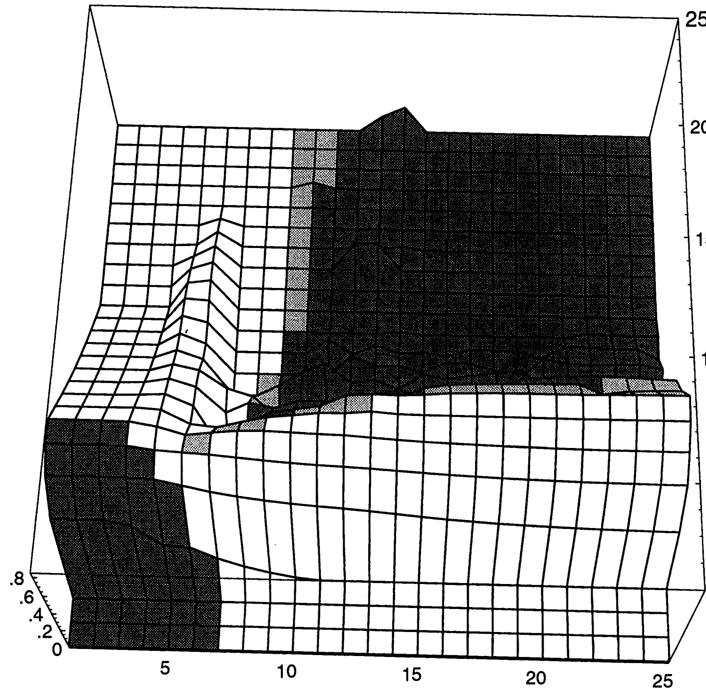

FIGURE 6

INVESTMENT (PROBABILITY OF SUCCESSFUL RESEARCH), $\omega_3 = 4$

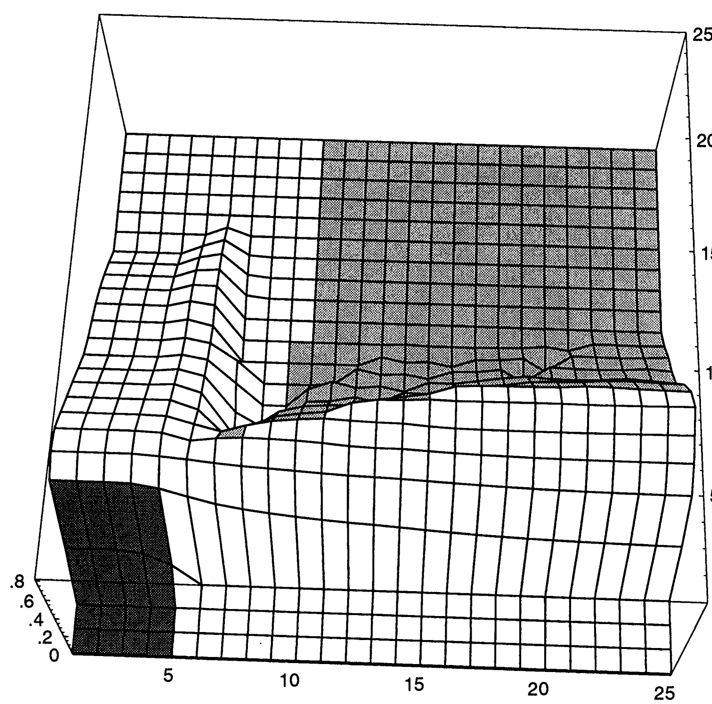

firms are colluding, and yet that collusion does not induce entry (this is the lightly shaded region in the figure). Any movement from these states downward will cause collusion to break down, while any movement upward will induce entry. When there is a third firm active at $\omega = 3$ (Figure 4), the band in which collusion can be sustained but entry is not induced becomes quite a bit larger. Moreover, if the third firm's ω increases to 4, collusion is maintained and entry disappears altogether (see Figure 5).

As noted, Figures 5 and 6 are a transform of the *derivatives* of the value function in Figures 3 and 4, and hence they provide more details on the shape of that function. It is clear that the firm invests heavily when successful research has a large impact on its likelihood of being in the collusive region. Moreover, when the ω's of the incumbents are both moderately high, and one firm pulls ahead of the other just enough to endanger the collusive agreement, the leader stops investing and lets the follower catch up.

Once firms are comfortably in the region where collusion can be sustained, investments go to zero and stay there. Apparently, the Nash bargaining solution implies that the benefits to increasing one's ω in this region are diffused so thoroughly among the incumbents that there is insufficient inducement to invest (this implies that tuples of incumbents with very high ω's will never be observed). Finally, there are also small "hills" in investment surrounding the entry region; hills that correspond to entry deterring and exit-inducing behavior.

Though not shown here, the value function and the investment policies for the model where firms are not allowed to collude are markedly different from those in Figures 3–4 and 5–6. First, the value functions with collusion are noticeably larger, especially in the collusive region (where it is a factor of two to three higher). Moreover, with no collusion there is no rise in the value function, and no corresponding increase in investment, around the borders of the collusive region. Investment of the leading firm increases when its competitor's ω approaches it from below, for it is precisely when two firms have similar ω's that competition is most intense. In contrast, when there was collusion, the leading firm decreases its investment when the second firm approaches it from below (for if the two firms' ω's become similar enough they can sustain collusion). Moreover, if in the noncollusive case we move out along a diagonal with both firms having similar ω's, the investments remain high for some time (instead of going fairly rapidly to zero, as in the collusive case). These observations should lead us to believe that an industry which can collude is likely to have quite a different distribution of both the number and the states of active firms than would an industry in which there are no collusive possibilities.

□ **Descriptive and welfare analysis.** We used these policies to simulate 100,000 periods of industry evolution, starting from an initial condition with no firms active. We then did the same for an industry in which collusion was not allowed. Statistics from these runs appear in Table 1.

As noted, the noncollusive industry is largely a "natural monopoly" (i.e., in 90% of the periods there is only a single firm). The industry is also quite "stable"; the average length of a run with the same firm monopolizing production is 68 periods (the last set of rows in Table 1). The quality of the monopolist's product does eventually fall low enough relative to advances in the outside alternative to induce entry. Entry is followed by a period of active investment competition among firms with low-quality products. The competitive phase is very unstable; there will be a third entrant if the investments of the first two are not successful, and it is quite likely that one or both of the firms that were initially competing, and some of the subsequent entrants, fail (this is why the modal life span is two periods and the median is three). This competitive phase typically ends with one firm pulling away from the others, and the others

TABLE 1 Market Structure

	Markov Perfect with Collusion	Markov Perfect no Collusion
Percentage of equilibria with n active firms		
$n = 1$	47.7	89.1
$n = 2$	30.7	10.0
Of these, with $I = 1$	(44.8)	Not relevant
$n = 3$	21.1	.8
Of these, with $I = 1$	(77.3)	Not relevant
$n = 4$	0.52	.1
Of these, with $I = 1$	(13.2)	Not relevant
$\bar{n}$	1.74	1.12
With entry	5.4	2.9
With exit	5.3	2.2
Statistics on life-span distribution		
Median	3.0	3.0
Mean	32.7	51.1
Average length of runs for given (n, I) couples[a]		
$(n, I) = (1, -)$	31.1	68.0
$(n, I) = (2, -)$	7.3	5.7
$(n, I) = (2, 1)$	3.3	Not relevant
$(n, I) = (2, 0)$	3.3	Not relevant
$(n, I) = (3, -)$	6.7	1.3
$(n, I) = (3, 0)$	1.9	Not relevant
$(n, I) = (3, 1)$	5.2	Not relevant
$(n, I) = (-, 1)$	7.3	Not relevant

[a] The (2, –) row represents runs with the same two firms active, (2, 1) runs with the same two firms colluding, etc. The (–, 1) row represents a run of continuous collusion (though not necessarily between the same firms).

finding it unprofitable to compete with the superior product developed by that firm exiting the industry.

When we "take away the antitrust authority" and simulate from the policies computed from the institutions that allow for collusion, we obtain an industry with *more* competition. The fraction of the periods in which there is a monopoly producer falls from 90% to 48%, and the length of the average run with the same firm monopolizing production is cut by more than half. The collusive industry is more likely to generate an entrant to compete with an incumbent monopolist.

This is a result of the fact that the possibility of future collusion induces the entrant to enter when the incumbent monopolist is at a higher ω (it enters with the incumbent at $\omega = 5$, in contrast to $\omega = 4$ when there is no collusion) and is in spite of the fact that an incumbent monopolist invests more when future collusion is possible. The combined effect of the monopolist facing earlier entry, and of it investing more when there is the possibility of future collusion, results in the quality distribution of the product produced by a monopolist when collusion is possible being better (in the stochastic dominance sense) then the quality of the product produced when collusion is not allowed.

The fact that when collusion is possible entry occurs with the incumbent monopolist at a higher ω also implies that initially the entrant will tend to be farther behind the incumbent monopolist. Moreover, since collusion cannot be sustained from the states realized after entry (see Figure 2), immediately after entry there will be full price and investment competition just as in the noncollusive case. Indeed, if the ω of the initial incumbent continues to fall we will see two successive entry periods and competition among three competitors with low ω's. This type of competition, however, does not last long (on average less than two periods), as one of the three firms typically falls behind its two competitors and exits. Thus with or without the possibility of future collusion, entry will be followed by an initial unstable period with a high likelihood of one or both of the initial firms failing and subsequent entry.

The major difference between the collusive and noncollusive cases is that when we allow for collusion, any initial success by one firm in the competitive phase does not invariably turn into that firm dominating the market. When collusion is allowed, the firm that fell behind still invests heavily, since it realizes that if it does get close to the higher-quality firm the two will be able to share collusive (instead of pure Nash) profits, while for the same reason the firm with the higher state is less averse to the smaller firm catching up and invests less heavily (see Figures 3 and 4). Consequently, the way out of the low-quality competition that follows entry when there is the possibility of future collusion is often for two firms to develop a fairly successful product, successful enough to enable them to collude. This also implies that the two-firm states we observe when collusion is allowed are generally higher-quality couples of states than the couples of two-firm states we observe when collusion is not allowed (just as was the case for the observed monopoly states).

Recall that the set of states at which the two firms collude and do not induce entry by a third firm is very narrow. Firms that are just inside the collusive region must invest in order to maintain states that are high enough to allow them to collude, but if the investments are too successful, the firms move to a couple of states that induce entry. It is for this reason that the length of the run with the same two firms colluding is typically quite short (on average 3.3 periods). As noted, when the two-colluding-firms states are high enough to induce entry, collusion will be maintained after entry. As a result, the average length of a run with collusion is a much larger, 7.3 periods.

Most of the entry from a state when two firms are active is entry to a collusive state (62% of it; as noted, the alternative is to a state where three small firms are competing). Consequently, we observe collusion in about 75% of the states when three firms are active. The three collusive firms typically invest enough to deter a fourth entrant and then let investment fall to almost zero (see Figures 5 and 6). Note, however, that when there is collusion, we do observe three firms active with two of them producing a high-quality product; when there is no collusion, the only times we see three active firms is when all their products are low quality.

So the industry with collusive possibilities generates quite different evolutionary patterns than the industry without collusion. The industry without collusion looks a lot like a natural monopoly that offers one relatively low-quality product. Occasionally it is challenged by a new product and a competition ensues. The competition continues, however, only so long as none of the competing firms has a run of successful investment. Once one firm gets ahead, it tends to attain a dominant position and we revert to monopoly.

In contrast, when collusion is allowed the monopolist tends to invest more and develop a higher-quality product. This is required to forestall a potential entrant who sees a future with more lucrative collusive possibilities. Despite its greater investment, the monopolist in the collusive industry is not nearly as successful in deterring entry, so that we see many more periods with more than one product offered. Further, since to sustain collusion the couple of firms must have qualities sufficiently high to make exit a remote possibility, the two-firm periods in the equilibria that allow for collusion have tuples of states that are quite high. Indeed, we frequently move to a tuple high enough to induce yet a third entrant. Consequently, not only are there more multiproduct periods in the collusive industry, but even conditioning on the number of incumbents, the quality of the products offered is typically higher when there are collusive possibilities than when there are none.

Though allowing for collusion tends to generate both more and higher-quality products for consumers to choose from, collusion need not make consumers better off: it depends on the prices that consumers have to pay for those products. We know that collusive prices are higher than noncollusive prices at a given tuple of states. So given the tuples of states, the consumers will prefer not to have collusion. The question of whether we wish to rid ourselves of collusive possibilities is then a question of whether the benefits from lower prices conditional on the states outweighs the losses from having both a smaller number and lower quality of products available.

Price information is provided in Table 2. That table has the average, over periods, of the sales-weighted average of the prices under the alternative institutional regimes. Recall that marginal cost is always 5, so one can read the markups directly off this table. Consumer surplus differs with both this markup and the ω's, so we require a separate calculation for that.

Perhaps the most striking fact from Table 2 is the difference between prices in collusive and noncollusive periods: conditional on the number of firms active, that difference is *over* 50%. Moreover, virtually all this difference occurs in the one-period transition between collusive and noncollusive regimes; there is on average a 50% increase in prices when collusion begins and a 50% fall in prices in periods when collusion breaks down. Thus the equilibrium generates price patterns that look very much like *price wars.*

Note that the collusive prices when they occur are (again on average) even higher than the monopoly price. Recall that to support collusion all colluding firms must have ω's that are quite high (see Figure 2). The single firm wants to deter entry, but it can do this at an ω that is typically lower than the ω's needed to sustain collusion. That is, the states of colluding firms are typically higher than the states of a monopolist; this explains why prices are typically higher when there are colluding firms than when there is a monopolist.

Indeed, the whole comparison between prices in collusive industries and prices in industries where collusion is prohibited is quite complicated. On average, prices in the noncollusive industry are slightly higher. This is because there are so many more monopoly periods when collusion is not allowed. But if we compare either just monopoly or just duopoly periods, the industry with collusion has higher prices. In monopoly periods, the higher prices are solely a reflection of the fact that the monopolist in the industry that allows for collusion is typically at a larger ω than a monopolist in an industry that does not allow collusion. When we look at periods when two goods

TABLE 2 Prices and Benefits

	Markov Perfect with Collusion	Markov Perfect no Collusion
Prices with n firms active[a]		
$n = 1$	11.0	10.9
$n = 2$	9.3	7.3
Of these, with $I = 1$	11.5	Not relevant
Of these, with $I = 0$	7.5	Not relevant
$n = 3$	10.4	6.4
Of these, with $I = 1$	11.5	Not relevant
Of these, with $I = 0$	6.6	Not relevant
Average	10.2	10.3
Average percentage price change		
When collusion begins	47.0	Not relevant
When collusion ends	−48.0	Not relevant
Consumer surplus[b]		
Mean	22.8	19.8
Standard deviation	6.3	7.2
Producer surplus[b]		
Mean	34.5	38.1
Standard deviation	11.1	10.7

[a] The prices are the average, over periods, of the sales-weighted average price.
[b] The consumer (producer) surplus meansures are the mean and standard deviation of the discounted sum of consumer (producer) benefits over a 100-year period averaged over 1,000 runs started at random draws from the ergodic distribution of states.

are offered, one of two different scenarios unfolds. If there is a price war, the prices that emerge from the industry with collusive possibilities are very similar to the average prices that occur when two firms are active in an industry in which collusion is not allowed (this occurs because price wars occur at relatively low ω states, and these are just about the only two-firm states observed when collusion isn't allowed). On the other hand, when collusion can be sustained, the prices of the colluding products are higher than just about anything we ever see in an environment in which collusion is not allowed.

This brings us to the consumer and producer surplus calculations (the last two sets of rows on Table 2). Producer surplus is the discounted sum of total profits minus total investment and entry costs plus any exit values, over a 100-year period. Consumer surplus is the discounted sum of consumer utility over the same period. Table 2 contains the means and the standard deviations of the figures over 1,000 separate samples from randomly drawn initial conditions (we ran one long run of length 100,000, and then broke the output up into 1,000 subsamples of 100 periods each).

There is quite a bit of variance in these figures over runs in both institutional environments; indeed, the standard error of each surplus measure (over samples) is more than twice as large as the differences between the average surplus with and without collusion. Note also that the average of the sum of consumer and producer surplus when we allow for collusion is virtually identical to the average of this sum when we do not allow for collusion (they are within one-half of a standard deviation of those averages). Thus a social planner whose decision were based on an unweighted sum of consumer and producer surplus would be indifferent between the environment that allowed for collusion and the one that did not.

The results from the consumer surplus calculations are rather striking. Recall that the consumer is typically offered a larger number of products when collusion is allowed, and the products offered are typically of higher quality. On the other hand, when there is collusion the consumer is offered those goods at relatively high prices (higher than the prices that would be generated at the same states if there were no collusive possibilities), and the low-quality (and low-priced) alternative products are typically not available. Still, consumer surplus is on average *significantly higher* when collusion is allowed (the difference between the two means is over six times its standard error). The fact that the consumer benefits are higher when we allow for collusion is entirely because of the difference in *dynamic* incentives (investment, and entry and exit). For any given state, prices will be higher and consumer surplus will be lower when collusion is not allowed. However, the distribution of states is so much more favorable to the consumer when collusion is allowed that its effect overcomes the negative impact of higher prices on consumer welfare. It follows that a social planner who gave more weight to consumer than to producer surplus would prefer a set of institutions that allowed for collusion to one that did not.

We should note here that if we had worked with a more realistic distribution of utility functions, in particular a distribution that allowed for differences among consumers in their sensitivity to prices due to differences in their incomes, then what we would have undoubtedly found out is that consumers with income greater than some amount would have preferred the collusive industry, but lower-income consumers would have preferred an environment where collusion is not allowed. That is, there are distributional consequences of the choice between allowing or not allowing for collusion, and we have abstracted from such considerations here. Nonetheless, the idea that collusion is necessarily bad for the consumer is simply wrong. A collusive industry has different dynamic incentives than an industry in which collusion is not allowed, and this fact may lead to a distribution of states that is so preferred by the consumers that they favor the collusive industry despite its higher prices.

It is interesting that producer surplus is (again on average) higher for the industry that cannot collude. As noted, when we see collusion it is because the expected discounted value of net cash flows to each incumbent is larger when the firms collude than when they do not (else collusion could not be sustained). On the other hand, there is free entry into the collusive industry, and this pushes the discounted value of the marginal firm toward the sunk cost of entry. With our parameters, the difference between the monopoly and duopoly values when there is no collusion is often quite large. As a result, even in situations where the monopolist is quite profitable there is very little incentive to enter. When there are collusive possibilities, the duopoly values are larger, often large enough to induce entry (indeed, we often see three active firms), and this dissipates the rents accruing to the existing incumbents.

Some final points: We have also calculated the equilibrium for the perfect cartel (a multiproduct monopolist who controls all entry, exit, investment, and pricing decisions to maximize the discounted sum of producer net cash flow) and for the social planner (who sets price equal to marginal cost and sets entry, exit, and investment decisions to maximize the discounted sum of consumer surplus minus the resultant investment costs). For this set of parameter values, the perfect cartel generates industry structures, and producer and consumer surpluses, that are very similar to those generated by the noncollusive Nash equilibria (this should not be surprising given the results in Table 1).

The social planner also has one firm active in the vast majority of the periods, but that firm develops its product to higher states than in the other institutional environments, and the planner's firm sells at marginal cost. Recall that our planner maximizes the sum of consumer and producer surplus, and *all* the planner's surplus is allocated to consumers. In spite of this, Table 2 indicates that if a consumer-oriented planner were given a choice between the noncollusive Nash solution, a solution which from the point of view of concentration measures looks very much the planner's solution, or the collusive Nash solution, a solution which typically involves marketing a larger number of higher-quality products at collusive prices, the planner would take the collusive outcome.

All these results are for an industry that has high-enough sunk costs relative to demand for there to be only a small number of firms when collusion is allowed, and that is often a monopoly when there are no collusive possibilities. On the other hand, we have traditionally worried about collusion in industries with small numbers of active firms. Often the firms being investigated for collusion couch their defense in terms of the implications of "destructive competition." If their argument is interpreted as saying that if they were forced into a more "competitive" (say Nash) pricing arrangement the industry would produce fewer (and possibly lower-quality) products, and that this would hurt consumers, our calculations clearly indicate that there are situations when the firms are right.

5. Extensions

■ We have two goals in this article. One is to present a particular model of collusive behavior and analyze its implications. The other is to provide a framework in which we can analyze the implications of collusion in a variety of dynamic settings. The framework has two important elements of dynamic interactions: (i) the strategic aspect, that is, we allow firms to condition their actions on the history of the interaction, and (ii) the structural aspect, that is, we allow for the existence of state variables that affect the profit function and evolve over time in a manner that is partially controlled by the firms' actions.

As we note in the article, if one is concerned only with these two aspects of dynamic interactions, our framework is quite flexible in the sense that we can change the more detailed assumptions to better suit a given institutional setting and then recompute and analyze the optimal policies. Further, the dynamic logic that underlies our questioning of the belief that collusion is necessarily bad for consumers carries over, in slightly different guises, to these alternative environments. In particular, by introducing dynamics we introduce "marginal" conditions with respect to entry and investment, and these conditions tend to limit the extent to which collusion can lead to a sustained increase in profitability, as well as transfer some of the benefits from collusion to consumers.[27]

[27] Consider, for example, a homogeneous goods industry that used a quantity-based Nash bargaining solution for the collusive outcome with a Cournot Nash in quantities solution for both the threat points and the reversion, and compare it to a standard Cournot solution when collusion is not allowed. The increase in

On the other hand, a large part of the literature on collusion has focussed on settings where asymmetric information is important. As an example, consider a model in which the collusive periods and the price war periods also differ in the information they provide on the firms' state variables. For simplicity, assume that during collusive periods firms do not observe the outcomes of the investments of their competitors and hence their competitors' current states, while during price wars all states are revealed. Price wars conclude with an agreement on collusive prices (or quantities) based on the knowledge revealed during the war. These prices determine the allocation of the collusive rents among the firms. The prices (or quantity quotas) remain in force as long as there are no firms that demand to change them. But, if one of the firms has good realizations of its investment efforts, then that firm may no longer be content with the prior allocation and might demand a change that would reflect its new superior position.

The new position of the incumbent with successful investments is not observed by rival firms. Moreover, the other incumbents cannot immediately accept each demand for a new division of the collusive rents, since this would induce firms to announce successful investments when there were none. Assume then that the demand for change is accompanied by a threat of price war, and that in equilibrium some of these demands are rejected and the price war ensues. While this model is beyond the scope of the current article, it could be developed in a way that provided a role for asymmetric information, as well as the structural and strategic aspects of collusion discussed above. The price wars would reflect disagreement on the allocation of the collusive profits, so the model and the resultant price wars have some similarities with bargaining models with incomplete information about player characteristics (models that generate delays and strikes as part of the equilibrium). The bargaining process is required because the environment changes periodically, and when a sufficiently large change occurs, the old collusive agreement is no longer acceptable to some of the agents. The price wars in such a model occur as part of a bargaining process in which one of the producers demands a renegotiation of the collusive agreement. This feature of price wars is, for example, central to Levenstein's (1997) description of collusive behavior among bromine producers between 1885 and 1914.

6. Concluding remark

■ Standard antitrust analysis of collusive behavior is static in that it conditions on "market structure" and discusses the implications of collusion on quantities or prices for the given structure. This ignores the impact of collusion on the incentives to launch new products, or to invest in existing ones. Clearly, for a given market structure (a given set of state variables) society is better off when the firms compete and do not collude. But whether or not we allow for collusion also impacts on the market structures that are likely to be developed, and once we take these dynamic effects into account it is not at all clear that an antitrust authority interested in maximizing social welfare should forbid collusion. Indeed, our example shows that if the worry is about collusive prices in a market that supports only a small number of firms, we may well be better off encouraging collusion than deterring it (especially if the incentives generated by

entry and investment incentives provided by the collusive possibilities imply that the periods in which collusion could not be supported would be accompanied by greater output, and lower prices, than we would typically see in an equilibrium that never allowed firms to collude. Alternatively, we could complicate our model by allowing for multiproduct firms. Then, depending on parameter values, either we would find the type of equilibrium described here or we would find that the greater threat of entry due to the collusive possibilities would induce the monopolist to deter entry by producing more products than the monopolist would produce were collusion impossible.

the collusive possibilities allow us to avert a monopoly situation). More generally, a blanket per se aversion to collusion seems to be clearly off the mark.

This article suggests, then, that there is a need to revise standard teaching and policy practice vis-à-vis collusion, and it sets forth the beginnings of a framework to enable this to be done. We say "beginnings" both because the framework, even if applicable to a given situation, is incomplete without filling in institutional detail on the industry of interest, and because the framework itself needs to be enriched in a nontrivial way to be appropriate for settings in which asymmetric information is an important component of how collusion works. So to get from these beginnings to a cogent analysis of any particular industry will require a fairly in-depth knowledge of institutional detail, parameter values, and theoretical and computational possibilities.

References

ABREU, D. "Extremal Equilibria of Oligopolistic Supergames." *Journal of Economic Theory,* Vol. 39 (1986), pp. 191–225.

———, PEARCE, D., AND STACCHETTI, E. "Optimal Cartel Equilibria with Imperfect Monitoring." *Journal of Economic Theory,* Vol. 39 (1986), pp. 251–269.

BAGWELL, K. AND STAIGER, R.W. "Collusion over the Business Cycle." *RAND Journal of Economics,* Vol. 28 (1997), pp. 82–106.

CAPLIN, A. AND NALEBUFF, B. "Aggregation and Imperfect Competition: On the Existence of Equilibrium." *Econometrica,* Vol. 59 (1991), pp. 25–59.

COMPTE, O., JENNY, F., AND REY, P. "Capacity Constraints, Mergers and Collusion." Mimeo, University of Tolouse, 1999.

DAVIDSON, C. AND DENECKERE, R. "Excess Capacity and Collusion." *International Economic Review,* Vol. 31 (1990), pp. 521–542.

ELLISON, G. "Theories of Cartel Stability and the Joint Executive Committee." *RAND Journal of Economics,* Vol. 25 (1994), pp. 37–57.

ERICSON, R. AND PAKES, A. "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies,* Vol. 62 (1995), pp. 53–82.

FERSHTMAN, C. AND MULLER, E. "Capital Investments and Price Agreements in Semicollusive Markets." *RAND Journal of Economics,* Vol. 17 (1986), pp. 214–226.

——— AND GANDAL, N. "Disadvantageous Semicollusion." *International Journal of Industrial Organization,* Vol. 12 (1994), pp. 141–154.

FRIEDMAN, J.W. "A Noncooperative Equilibrium for Supergames." *Review of Economic Studies,* Vol. 38 (1971), pp. 1–12.

——— AND THISSE, J.-F. "Partial Collusion Fosters Minimum Product Differentiation." *RAND Journal of Economics,* Vol. 24 (1993), pp. 631–645.

FUDENBERG, D., LEVINE, D.I., AND MASKIN, E. "The Folk Theorem with Imperfect Public Information." *Econometrica,* Vol. 62 (1994), pp. 997–1039.

GREEN, E.J. AND PORTER, R.H. "Noncooperative Collusion Under Imperfect Price Competition." *Econometrica,* Vol. 52 (1984), pp. 87–100.

HALTIWANGER, J. AND HARRINGTON, J.E. "The Impact of Cyclical Demand Movements on Collusive Behavior." *RAND Journal of Economics,* Vol. 22 (1991), pp. 89–106.

HARRINGTON, J.E. "Collusion Among Asymmetric Firms: The Case of Different Discount Factors." *International Journal of Industrial Organization,* Vol. 7 (1989), pp. 289–307.

JUDD, K.L. *Numerical Methods in Economics.* Cambridge, Mass.: MIT Press, 1998.

KANDORI, M. "Correlated Demand Shocks and Price Wars During Booms." *Review of Economic Studies,* Vol. 58 (1991), pp. 171–180.

LEVENSTEIN, M. C. "Price Wars and the Stability of Collusion: A Study of the Pre-World War I Bromine Industry." *Journal of Industrial Economics,* Vol. 45 (1997), pp. 117–147.

MASKIN, E. AND TIROLE, J. "A Theory of Dynamic Oligopoly: I and II." *Econometrica,* Vol. 56 (1988), pp. 549–569 and 571–599.

——— AND ———. "Markov Perfect Equilibrium, I: Observable Actions." Harvard Institute of Economic Research Working Paper no. 1799, 1997.

NASH, J.F. "The Bargaining Problem." *Econometrica,* Vol. 18 (1950), pp. 155–162.

PAKES, A. "Dynamic Structural Models: Problems and Prospects." In C. Sims, ed., *Advances in Econometrics: Sixth World Congress,* Vol. II. New York: Cambridge University Press, 1994.

——— AND McGUIRE, P. "Computing Markov-Perfect Nash Equilibrium: Numerical Implications of a Dynamic Differentiated Product Model." *RAND Journal of Economics,* Vol. 25 (1994), pp. 555–589.

———, GOWRISANKARAN, G., AND McGUIRE, P. "Implementing the Pakes-McGuire Algorithm for Computing Markov Perfect Equilibria." Mimeo, Yale University, 1995.

PORTER, R.H. "Optimal Cartel Trigger Price Strategies." *Journal of Economic Theory,* Vol. 29 (1983a), pp. 313–338.

———. "A Study of Cartel Stability: The Joint Executive Committee, 1880–1886." *Bell Journal of Economics,* Vol. 14 (1983b), pp. 301–314.

ROTEMBERG, J.J. AND SALONER, G. "A Supergame-Theoretic Model of Price Wars during Booms." *American Economic Review,* Vol. 76 (1986), pp. 390–407.

RUBINSTEIN, A. "Equilibrium in Supergames with the Overtaking Criterion." *Journal of Economic Theory,* Vol. 21 (1979), pp. 1–9.

SCHERER, F. *Industry Market Structure and Economic Performance.* Chicago: Rand McNally 1980.

SCHMALENSEE, R. "Competitive Advantage and Collusive Optima." *International Journal of Industrial Organization,* Vol. 5 (1987), pp. 351–367.

STAIGER, R.W. AND WOLAK, F.A. "Collusive Pricing with Capacity Constraint in the Presence of Demand Uncertainty." *RAND Journal of Economics,* Vol. 23 (1992), pp. 203–220.

STARR, A.W. AND HO, Y.C. "Nonzero-Sum Differential Games." *Journal of Optimization Theory and Application,* Vol. 3 (1969), pp. 1984–2008.

STIGLER, G.J. "A Theory of Oligopoly." *Journal of Political Economy,* Vol. 72 (1964), pp. 44–61.

SYMEONIDIS, G. "Are Cartel Laws Bad for Business? Evidence from the UK." Mimeo, University of Essex, 1999.