A DYNAMIC OLIGOPOLY WITH
COLLUSION AND PRICE WARS

Chaim Fershtman
Ariel Pakes

A Dynamic Oligopoly with Collusion and Price Wars
Chaim Fershtman and Ariel Pakes
NBER Working Paper No. 6936
February 1999
JEL No. C72, C73, D43, L4

## ABSTRACT

Most of the theoretical work on collusion and price wars assumes identical firms and an unchanging environment, assumptions which are at odds with what we know about most industries. Further that literature focuses on the impact of collusion on prices. Whether an industry can support collusion also effects investment incentives and hence the variety, cost, and quality of the products marketed.

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 either one of the firms does not keep up with the advances of its competitors, or a "low quality" entrant enters. In either case there will be an active firm that is quite likely to exit after it deviates, but if one of the competitors is near an exit state the other incumbent(s) has an incentive to price predatorily (that is to deviate themselves).

We use numerical analysis to compare an institutional structure that allows for collusion to one which does not (perhaps because of an active antitrust authority). Price paths clearly differ in the two environments; in particular only the collusive industry generates *price wars*. The collusive industry offers both more and higher quality products to consumers, albeit often at a higher price. The positive effect of collusion on the variety and quality of products marketed more than compensates consumers for the negative effect of collusive prices, so that *consumer surplus is larger in the collusive environment.*

Chaim Fershtman
Berglas School of Economics
Tel Aviv University
Tel Aviv 69978
Israel
fersh@tccsg.tau.ac.il

Ariel Pakes
Department of Economics
Yale University
PO Box 20864
New Haven, CT 06520-8264
and NBER
ariel@econ.yale.edu

# 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 loathe to assume that the different firms in a market have the same policy cum profitability options. This is particularly unfortunate since there is a very limited range of pricing models currently available 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 data sets.

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 effects 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.

Perhaps not surprisingly, much of the literature on collusion developed with an explanation for the price war phenomena in mind. It was Stigler (1964) who first pointed out the possibility that price wars are 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 early 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)).

---

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

The early articles did make it clear that to understand collusion we need to understand 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, henceforth GP) were the first to obtain price wars as part of the equilibrium behavior. GP considered a repeated oligopolistic interaction with imperfect monitoring in which demand is subject to random unobservable shocks and the firms' outputs are unobservable by their rivals. When a low price is observed, firms do not know if this low price is the consequence of a deviation from collusive pricing by one of their competitors or if there was, in fact, a low realization of the demand function. Green and Porter (1984) 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) that 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.[2] A second model with equilibrium price wars was introduced by Rotemberg and Saloner (1986) that modeled oligopolistic collusion with time varying demand. In their setup a period with high demand, or a "boom", generates a greater temptation to deviate from the collusive agreement and hence countercyclical pricing. [3]

We also allow firms to condition their actions on the history of the interaction, thereby keeping the focus of our discussion on the collusive possibilities generated by dynamic equilibrium. However we allow, in addition, for investments to effect 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

---

[2]See also Porter (1983a,b) and Ellison (1994).

[3]For further analysis of this model see Bagwell and Staiger (1997), Staiger and Wolak (1992), Kandori (1991) and Haltiwanger and Harrington (1991). Note that the first model predicts that we observe price war during periods with low demand while the second predicts such behavior during periods of high demand. For more on the comparison of the two models and an empirical assessment of their applicability to the 1880 railroad cartel (the Joint Executive Commitee), see Ellison (1994).

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 lead to.

Moreover since the outcome of the investment and entry processes is stochastic, over time the different firms will find themselves in different states with 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 that are likely to be achieved.

We begin by adapting the framework presented in Ericson and Pakes (1995; henceforth EP) 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 which determines those profits, and an investment game which 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 they play noncooperatively with respect to investment.[4] As in EP equilibrium is Markov Perfect, but now strategies are allowed to depend on past pricing behavior, as well as on the "payoff relevant" states studied in Maskin and Tirole (1988) and used in EP.[5]

The model allows for exit from and entry into the industry. Exit rewards firms with a fixed sell-off value of its equipments. When the future prospect of staying in the industry is below the sell-off amount, the firm exits. As in EP, the specification for the evolution of the state variables, together with equilibrium decision making, insures that each firm exits the industry in a

---

[4]Implicitly we are assuming prices are easy to observe, while investment, at least the exploratory investment we focus on, 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. For further analysis of semicollusive markets see Fershtman and Gandal (1994) and Friedman and Thisse (1993).

[5]Note, however, that information remains symmetric, as in Maskin and Tirole, 1995; for a discussion of extensions see our section 5.

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, and second the fact that one or more firms might exit provides an incentive for the continuing incumbent(s) to prefer non-collusive (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 computation algorithm developed by Pakes and McGuire (1994); (henceforth PM), to allow for the possibility of collusive behavior.[7] To allow us to isolate the aspects of our results that are due to collusion we compare the numerical results from an institutional structure that allows for collusion to an industry with the same cost and demand parameters but with no collusive possibilities (perhaps because of an active antitrust authority).

Price fixing is prohibited by Section 1 of the Sherman Act. Over the years courts have taken a clear position with regard to collusive behavior such as price fixing. Collusive practice is per se illegal which means that it is sufficient to prove the existence of such a practice without the need to prove its harmful effects. This position of the antitrust law towards price fixing reflects the standard economic analysis which claims that price fixing indeed reduces efficiency. Clearly, at any particular period, prohibiting price fixing will result in a larger consumers surplus for that particular period but

---

[6]Of course the exit of one competitor might be followed by the entry of another, but there are sunk costs and time required to enter, thus insuring that the continuing incumbent will be temporarily better off. We should note that we assume that at each market structure only a small number of possible price vectors will be calculated and hence are in the choice set, and that these must be calculable from the current profit function. In particular firms will chose between (relatively easy to calculate) collusive prices, non-collusive prices, and deviating prices. Thus we do not consider the problem of determining which prices could be supported at each market structure, and hence we do not consider the makeup of the optimal cartel. The determination of the set of prices that could be supported is a difficult problem, so in order to base our model on them we would have to require firms to perform and agree on a far more complex set of calculations (see the discussion below), but were it possible to do so, we would expect that the use of optimal cartel prices to ameliorate the rather sharp movements in prices we report below.

[7]Though numerical analysis can be reasonably realistic, it has to be based on very detailed assumptions. Our hope is that eventually researchers will be able to choose these assumptions on the basis of empirical results on the industry of interest.

the possibility of collusion and the type of the collusive equilibrium affect the firms' investments as well as their exit and entry decision. In such a case the market structure is endogenously determined and is affected by the collusive process. Comparing the equilibrium with and without collusive possibility indicates that the collusive industry is less concentrated 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 collusive prices, so that consumer surplus is larger in the collusive environment. Our analysis therefore indicates that the acceptable idea that collusion is necessarily bad for consumers is false. It is possible that in the long run consumers will be better off when collusion is not prohibited. The paper, therefore, suggests the need to revise the standard anti-trust policy towards collusion. While clearly our result is not general enough to reverse the recommendation of the antitrust policy with respect to price fixing, it nevertheless emphasis the need to reconsider the per se illegality of such a practice.

## 2 The Model.

We adapt the framework presented in Ericson and Pakes (1995; henceforth EP), and the algorithm for computing it presented in Pakes and McGuire (1994; henceforth PM), to allow for collusion. In each period there are $n_t$ incumbent firms which differ in their physical characteristics, in say $\omega_{i,t}$. $\omega_{i,t}$ is a characteristic of either the product produced by the firm or its cost function, and it evolves over time with the outcomes of an investment process. Positive outcomes lead to states in which the firm can make more 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 on 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 prices, say $p_{i,t}$, as a function of history. These prices, together with the firms' state variables determine

5

the profits of each active firm, say $\pi(i, t) = \pi(\omega_{i,t}, p_{i,t}, \omega_{-i,t}, p_{-i,t})$. Investments are directed at improving the firm's "physical" state, their $\omega$ value. The period concludes with the realizations of the stochastic outcomes of the investment and entry decisions [8].

## 2.1 Physical States, Investment, and Entry and Exit.

As in EP(1995) we assume that $\omega_{i,t}$ takes on values in the positive integers, $\omega_{i,t} \in \Omega \subset \mathcal{Z}^+$, so $\omega_t = (\omega_{1,t}, \ldots, \omega_{n_t,t}) \in \Omega^{n_t} \subset (\mathcal{Z}^+)^{n_t}$. $\omega_{i,t}$ evolves over time with the outcomes of the firm's investment process, say $\eta_{i,t}$, and an industry specific exogenous process that affects the profit opportunities facing the industry in a given period, say $\nu_t$. $\nu_t$ captures the effects of factor prices and improvements in competition from products outside the industry; factors which generate positive correlation among the profits of our competing firms. Thus

$$\omega_{i,t+1} = \omega_{i,t} + \eta_{i,t+1} - \nu_{t+1} \qquad (1)$$

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

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

which is assumed stochastically increasing in x. The distribution of $\nu$ is given exogenously.

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

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 there is at most one entrant in every period, and indicate whether entry occurs by the indicator function $\chi^e = \{0, 1\}$, $\chi^e = 1$ indicating entry [9].

---

[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.

[9]Note that the subset of $\omega$ at which the entrant enters is independent of the general

## 2.2 Profits Conditional on Prices.

The version of the PM(1994) algorithm currently available computes equilibria for different types of markets as determined by the profit function used in the calculations[10]; but PM 1994 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, and to make it easy to compare our results to the noncollusive results available we use the demand and cost structure used in PM (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 $j = 1, \ldots, J$ goods in the market being studied, or chooses to spend all of its income on the "outside alternative" (good 0). The consumer who chooses good $j$ obtains utility $U_{ij} = g(\omega_j) + (y_i - p_j) + \epsilon_{ij}$, where $\omega_j$ 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 $\epsilon$ 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 which maximizes its utility.

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 $lim_{\omega \to \infty} g(\omega) < K$.[11]

---

progress of the industry; i.e. of the realizations of $\eta$ and $\nu$. Thus entrants improve with the improvement of knowledge in the industry. If this did not occur entry would eventually go to zero and stay their.

[10]A program to implement the algorithm can be accessed by FTP, , and is described in Pakes, Gowrisankaran, and McGuire, 1995. In addition to the differentiated product market described below, this program lets the user specify one of two homogeneous goods industries. In one costs are constant until a capacity constraint which responds to investment, and in the other firms invest to decrease marginal cost. The program has three modules. In the first the user chooses the form of the profit functions and sets demand and cost parameters. The program then calculates profits at all states. The second module takes these profits as input and calculates Markov Perfect investment, entry and exit policies, and the third takes the policies as inputs and then simulates market outcomes from a user specified initial condition. For comparisons the program also contains subroutines for obtaining similar output for the social planner's and perfect colluder's problem. There is a readme file to start you off.

[11]More exactly $exp(g(\omega)) = exp(\omega * [2 - exp - (\omega - \omega*)]$ for $\omega > \omega*$ so the limit of $g(\cdot)$

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 $\epsilon$), say $\sigma_j \equiv \sigma(\omega_j, \omega_{-j}, p_j, p_{-j})$ is given by

$$\sigma_j = \frac{exp[g(\omega_j) - p_j]}{1 + \Sigma_{q=1}^{J} exp[g(\omega_q) - p_q]} \tag{2}$$

Consequently if marginal cost is the constant $c$ then, conditional on any set of prices, the profits of firm $j$, say $\pi_j$, are

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

## 2.3 Equilibrium.

A subgame perfect equilibrium for the above game consists of a collection of strategies that constitute a Nash Equilibria 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. Thus we allow strategies to depend upon both the "payoff relevant" physical states, the $\omega_t$ used in Maskin and Tirole (1988 and 1995), and on a set of indicator functions which 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 $\alpha_i \in \{0, 1\}$, to indicate which, if any, of the existing firms have deviated in the past (in which case $\alpha_i = 1$). So the state of the system in period $t$ will be characterized by the couple $(\omega_t, \alpha_t) \in \Omega^{n_t} \times 2^{n_t}$. In the model we compute $\alpha_{i,t}$ evolves in a simple way: when a new firm enters the industry its $\alpha_i = 0$, and it only becomes one if the firm deviates. If the firm does deviate its "$\alpha$" stays at one for the remainder of its life.[12]

---

is $2exp(g(\omega*))$.

[12]One can also extend the setup and consider the case in which the indicator function stays at 1 only for some finite number of periods. Such a modification would enable us to consider equilibria in which punishment is limited for finite number of periods. However, such an extension would add to our computation complexity as the state space would be much larger.

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

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

## 2.4 Pricing and Collusive Behavior.

Since we are allowing for asymmetric firms even simple pricing rules generate prices which 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" that the regulatory authorities might trace to them). Thus, in the spirit of Schmalensee (1987), we keep matters relatively simple. In particular we limit the pricing possibilities to a collusive, a deviant, and a non-collusive price, and require that each of these prices can be computed from knowledge of the form of the current profit function (equation 3). We do, however, insure that collusive prices will only be observed 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 that the firm choose do not affect the evolution of the physical state, $\omega$. This implies that one equilibrium of the pricing game is the Nash equilibrium to the static (one shot) pricing game. The static Nash equilibrium insures that each firm's price maximizes its current profits (given by equation 3) given its competitor's prices and the characteristics of all products. As a result these prices satisfy the vector of

---

[13]Clearly such a Markov Perfect equilibrium is also a subgame perfect equilibrium in the game with the general history dependent strategy space.

9

first order conditions

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

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 $\omega$ tuple will be denoted by $\pi^N(\omega_j, \omega_{-j})$ (they are obtained by substituting $p^N$ for $p$ in equation (3)).

The equilibrium obtained when these pricing strategies are always followed is the equilibrium that PM(1994) analyze. We will compare it to an arrangement which 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 $p^C(\omega_i, \omega_{-i})$ and the price a firm would charge if it were to deviate from the collusive agreement. Collusive profits are obtained by substituting the collusive prices into the profit function in (3) and will be denoted by $\pi^C(\omega_i, \omega_{-i})$

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

$$max_{[p^D]}\pi(\omega_j, \omega_{-j}, p^D, p_{-j}^C), \tag{5}$$

with deviant profits given by $\pi(\omega_j, \omega_{-j}, p_j^D, p_{-j}^C)$ (with $\pi(\cdot)$ from (3)). Note that since *any* deviation is followed by a punishment mode by which the industry is playing the Nash Equilibrium prices $p^N(\omega)$ , the deviation is indeed the optimal deviation strategy.

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 to price to maximize 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 which 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. Since we assume that when the collusive prices cannot be supported prices revert to $p^N(\cdot)$ and profits to $\pi^N(\cdot)$, we take the threat point

for the bargaining game to be the one shot Nash equilibrium payoffs, $\pi^N(\cdot)$. Thus the collusive pricing vector, say $p^C(\cdot)$, is the solution to

$$\max_{p_1,\ldots,p_{n_t}} \Pi_{j=1}^{n_t} [\pi(\omega_j, \omega_{-j}, p_j, p_{-j}) - \pi^N(\omega_j, \omega_{-j})]_+, \qquad (6)$$

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 $\omega$ tuple are obtained by substituting $p^C$ into the profit equation in (3) and will be denoted by $\pi^C(\omega_j, \omega_{-j})$.

This completes the outline of the primitives 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. We come back to this point in the last section of our paper where we note several of the modifications which can be made to our model. Indeed perhaps the biggest benefit of our framework is that it can be perturbed to mimic many different institutions quite easily. Thus one way of reading this paper is as a modeling framework that can help map knowledge of the industry's institutions into the industry's collusive possibilities.

Our model does, however, have several feature which are more realistic than most models used in the past to analyze collusion. In particular we allow for heterogenous firms whose states evolve over time according to the outcomes of an investment process, and 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 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, 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 multi-product monopoly).

# 3 Computing the Equilibrium

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

As in that paper we assume that $\omega$ resides on the integers and invoke arguments analogous to those in EP(1995) to show that in equilibrium:

- We only observe $\omega$ values in a finite set, say $[1, \ldots, \overline{\omega}]$, and

- 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 $j^{th}$ firm is not active, then the observed $\omega$ tuples take values in $\Omega^{\mathcal{N}}$ where $\Omega = [0, 1, \ldots, \overline{\omega}]$.

Recall that in our model behavior depends both on the $\omega$ 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 $(\omega, \alpha) \subset \Omega^{\mathcal{N}} \times 2^{\mathcal{N}}$.

We also use PM's adaptation of the EP model for the transitions of $\omega$ for the firms who remain active. Since $\omega_j$ 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 $\omega$ 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[14]. If we take the model's decision period to be so small that only unit increments in both the outside alternative and in the quality of the firm's product can occur, let $p(x)$ be the probability that the firm's product improves and $\delta$ be the probability that the outside alternative improves we have: $\omega_{j,t+1} = \omega_{jt} + 1$ with probability $p(x)[1 - \delta]$, and $\omega_{j,t+1} = \omega_{jt} - 1$ with probability $[1 - p(x)]\delta$[15].

Firms which invest more have a larger chance of improving their $\omega$, so we require $\partial p(x)/\partial x \geq 0$. Moreover we assume both that $p(0) = 0$ (so that

---

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

[15]Note that we can allow for an arbitrary number of "model periods" for each time period in a given data set (though $\beta$ would have to be adjusted accordingly). As a result this specification can generate quite complicated patterns of $\omega$ transitions over the periods in a given data set.

a firm cannot improve its $\omega$ 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) which 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)). We now describe the Bellman equation which defines the (fixed point which determines the) value functions. 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.

## 3.1 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") determines when the industry can support collusion, and hence must be 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

$$S \equiv \{(\omega, \alpha) | \alpha \in 2^{\mathcal{N}} and \ \Sigma\alpha_j \in \{0, 1\}, \ \omega \in \Omega^{\mathcal{N}}\}.$$

We first present the Bellman equation, and then explain it verbally. For each $(\omega, \alpha) \in S$, the value of the $j^{th}$ firm satisfies

$$V(\omega_j; \omega_{-j}, \alpha) = max\{\phi, \ \pi(\omega_j; \omega_{-j}, \alpha) \tag{7}$$

$$+max_{x \geq 0}[-x + \beta\Sigma_{\omega'}V(\omega'_j; \omega'_{-j}, \alpha')p(\omega'_j|\omega_j, x)p(\omega'_{-j}|\omega, \alpha)\},$$

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

$$I(\omega, \alpha) = 1 \ \textit{if and only if}$$

- $\alpha = 0$, and

13

- for all $j$

$$\pi_j^D + max_{x_j}[-x_j + \beta\Sigma_{\omega'}V(\omega'_j; \omega'_{-j}, \alpha')p(\omega'_j|\omega_j, x_j)p(\omega'_{-j}|\omega, x, \alpha_j = 1, \alpha_{-j} = 0)]$$

$$< \pi_j^C + max_{x_j}[-x_j + \beta\Sigma_{\omega'}V(\omega'_j, \omega'_{-j}, \alpha')p(\omega'_j|\omega_j, x_j)p(\omega'_{-j}|\omega, x, \alpha = 0)].$$

The first max operator compares the exit value of the firm ($\phi$) to its continuation value. If $\phi$ 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 as 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 profits . If $\alpha = 0$ then the firms only collude 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, of $\omega'_{-j}$, 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 $\omega_e$ and their probabilities, $p_e$, are exogenously specified). The entrant enters if it is profitable to do so; that is if

$$\beta\Sigma_{\omega'}V(\omega_e; \omega'_{-e}, \alpha')p_e(\omega_e)p(\omega'_{-e}|\omega, \alpha) - x_e > 0. \tag{8}$$

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

Note also that $\alpha'$ is determined by $\omega'$ and $\alpha$. If $\alpha = 0$ then, since neither incumbents nor entrants deviate in equilibrium, $\alpha' = 0$. If $(\alpha_i, \alpha_{-i}) = (1, 0)$ for some $i$ then in states where the $i^{th}$ firm remains active, $\alpha'$ has a one in the $i^{th}$ slot and zero elsewhere, while if the $i^{th}$ 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

14

who exit are assumed not to make profits in the current period. Since neither do entrants, profits are calculated from the $\omega$ tuple of the incumbents who 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.

## 3.2 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)). [16]

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 $(\overline{\omega}, \mathcal{N})$.

In memory we hold

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

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

    - $\pi^N(\cdot)$,

- $\pi^D(\cdot), \pi^C(\cdot)$ 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 predeter-

---

[16]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. The Pakes and McGuire paper also formally introduces the operators which define the algorithm (a step that seemed unnecessary here given that it has already been done once).

mined order performing the following sequence of calculations [17]. 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 $x's$ appearing in equation (8), 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 $\alpha \neq 0$ (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 $x's$ that appear in the conditions determining whether a collusive equilibrium can be supported, and then check whether it can (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

$$max_{x \geq 0}[-x + \beta \Sigma_{\omega'} V^k(\omega_j'; \omega_{-j}', \alpha') p(\omega_j'|\omega_j, x) p(\omega_{-j}'|\omega_{-j}, x_{-j}^k, \chi_e^{k+1}),$$

where $p(\omega_{-j}'|\omega_{-j}, x_{-j}^k, \chi_e^{k+1})$ is notation for the probability of $\omega_{-j}'$ 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

$$V_j^{k+1}(\cdot) = \pi_j[\cdot, I^{k+1}(\cdot)] - x_j^{k+1} +$$

$$\beta \Sigma_{\omega'} V(\omega_j'; \omega_{-j}', \alpha') p(\omega_j'|\omega_j, x_j^{k+1}) p(\omega_{-j}'|\omega_{-j}, x_{-j}^k, \chi_e^{k+1}).$$

If this value is greater then the sell off value of the firm $(\phi)$, we set $\chi_j^{k+1}$, the indicator function for whether the firm exits, to zero, otherwise this indicator function is set to one.

---

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

16

If none of the incumbents exit, we copy the continuation values (i.e. the $V_j^{k+1}(\cdot)$) given above into memory, thus completing the sequence of calculations for the current point for this iteration. If one or more of the incumbent's have continuation values less than $\phi$, we set those firm's $V_j^{k+1}$ to $\phi$ and copy over all policies (entry, collusion, investment, and exit policies) from the $(\omega, \alpha)$ point obtained from the current $(\omega, \alpha)$ by resetting the $(\omega_j, \alpha_j)$ couples 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 of the investment policies calculated at successive iterations are below some critical value [18]. 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.

Note that we use last iterations 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

- check the entry condition

- check the collusion condition

- solve a set of single agent optimization problems (for the investment of incumbents), and

- 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[19]. Of course at the fixed point the current

---

[18]In the calculations below our convergence criterion was that the $R^2$ between successive iteration's values $(x's)$ were within $10^{-5}$ of one *and* the difference between successive iterations means were also less than $10^{-5}$.

[19]This would have essentially turned the problem into a nested fixed point problem; an inner fixed point would have had to be calculated at each point over which we cycle for the outer fixed point.

iteration's policies are identical to those of the last iteration, so the solution for the policies *is* this simultaneous solution.

We now come back to the problem of obtaining $S$. As in Pakes and McGuire, 1994, the 0 and $\bar{\omega}$ which determine $\Omega$ 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 [20]. $\mathcal{N}$ (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 $\mathcal{N} = \mathcal{N}^*$. [21]

# 4   Numerical Results.

The precise parameter values used to obtain numerical results are similar to those in PM(1994) (a paper which 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 $(\phi)$ stays at .1, so that the unrecoverable entry costs goes up from .1 to .7. Entry costs are now about $(1/6)^{th}$ of total production costs within a period compared to about $(1/125)^{th}$ in PM . As in PM the entrant who pays the sunk costs enters in the following period at $\omega = 4$ if the outside alternative does not move up in the interim, and at an $\omega = 3$ if

---

[20]It can be shown that an $\omega$ 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 (which follows from the boundedness of the profit function and the fact that $\beta < 1$) implies that there exist an $\omega$ above which no firm invests. Since without investment, a firm cannot improve its $\omega$, this puts an upper bound on the $\omega's$ that will ever be observed. Numerically we have found that the $\omega$ at which the monopolist stops investing is higher than the $\omega$ which stops 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 $\omega = \bar{\omega}$ actually wants to invest.

[21]We note that as in Pakes and McGuire, 1994, 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 PM. 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.

it does move up [22].

PM's(1994) base case had the maximum number of firms ever active $\mathcal{N} = 6$, and 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 rate of returns 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 much less frequent. Thus, as we shall see, we are dealing with an industry whose primitives are associated with a smaller number of longer lived firms, and less entry and exit. Thus we introduce the possibility of collusion in a fairly stable environment, one where one might think collusion could be reasonably successful[23]

We begin with a description of the policy output (of where collusion and entry occurs, and of the shape of the investment and value functions), and then provide descriptive and normative data from simulating the markets evolution over time.

## 4.1 Policy Output.

**Collusive States.**

Figure 1 looks at three firm states, first asks which of the three firms will continue and which will exit, and then 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 exits when its $\omega$ is two or three depending on the states of its competitors.

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 $\omega's$, and collusion breaks down because of a low $\omega$ of the

---

[22]The other parameters are set at: $\beta = .9$, $\delta = .65$ ($\delta$ is the probability that the outside alternative moves up), and $c = 5$ (recall that $c$ is marginal cost and it is denominated in the same units as $M$).

[23]There are other advantages of dealing with a smaller industry. One is that when there are one to three firms active (instead of PM's three to five firms active) it will be much easier for us to graph the results. Also our use of a smaller market decreases both the memory and the cpu time requirements for the computations.

third firm. However when two of the three firms have a high $\omega$ and the third firm does not exit (its $\omega \geq 4$), the three firms always collude.

That is, we only observe breakdowns of collusion among three continuing firms when *two of the three* firms have $\omega's$ that are near (but still above) the exit state. Just how low depends on how high the $\omega$ of the leading firm. The higher the leading firm's $\omega$ the smaller the area in which collusion can be supported. Thus collusion can be supported at $(8, 6, 6)$, but cannot be supported at any $(\omega \geq 9, 6, 6)$, etc. As we shall see the higher the $\omega$ 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 $\omega$, then they will collude as long as that $\omega \geq 5$. On the other hand when their $\omega's$ differ, collusion can break down even when both values of $\omega$ are much higher than that. Thus collusion cannot be supported at $\omega$ couples equal to either $(5, \omega \geq 6)$ or $(8, \omega \geq 12)$ (we see the collusive region when two firms are active more clearly below)[24].

There are two reasons to think that it might be harder to support collusion when 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 (or firms) realize 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, and since the smaller firm is more likely to exit if there is no collusion, the larger firm has an incentive to deviate in order to increase the likelihood of the smaller firm exiting[25].

---

[24]Though $\mathcal{N} = 4$, the results below indicate that we will rarely observe periods with $n = 4$, so we ignore those states here.

[25]Since the spot market equilibrium is Nash in prices when the firms are not colluding, none of the firms ever charge a price below marginal cost. Thus the use of the phrase "predatory" here 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.

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 which were 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 firm's preferred to deviate it was always the largest and the smallest firm.

This is an industry which is extremely profitable to a monopolist with an even modestly large $\omega$. The relatively large entry costs imply that a monopolist with a moderately large $\omega$ 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. 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[26].

---

[26]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 ( i.e. the value function when a single firm is active is higher if that firm is a deviant). This is because an incumbent who is a deviant does better at deterring future entry. This point to several other modeling possibilities. For example 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 "accomodating " the entrant (on the other hand such behavior might run the risk of trigerring a predatory pricing investigation). Alternatively we could assume that punishment phases last only a finite number of periods, and/or terminates 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). These and other variants of the model should not be too difficult to compute; indeed perhaps the biggest advantage of the numerical framework provided in this paper is that one can adapt it to study different institutional arrangements. Our purpose here is not to illustrate what "could" happen under different institutional arrangements, though we would find either a theoretical or a numerical analysis of this issue interesting. Rather we are illustrating the use of a tool

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 $\omega_e$) 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 produces a rather complex set of entry regions.

### Entry States.

Figure 2 provides a 3-dimensional plot of the states at which entry occurs (the shaded area in the figure). 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 have two levels). The plateaus correspond to entry states in which multiple firms are active and at least two of their $\omega's$ are quite large. The hill at the origin corresponds to entry states in which all active firms have very low values of $\omega$.

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 $\omega$, 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. However when there are three firms active then the inducement generated by collusive profits will only generate entry when the third incumbent has an $\omega = 3$ (the third incumbent would exit were its $\omega$ 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 the incumbents have *relatively low* $\omega$. In this case the entrant enters despite the fact that it will not earn collusive profits because it has a reasonable probability of either becoming a large dominant player in the future, or of becoming one of a small set of profitable future colluders (see below).

Entry *does not* occur when two firms are active when those firms have moderately high $\omega's$; i.e. $\omega's$ which are larger than the likely post entry

---

which can be adapted to a variety of situations.

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 $\omega$ is at least moderately large ($\geq 5$). This 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 (though the latter are not pictured here). The plateaus simply *disappear* when collusion is not allowed. Further without collusion when two firms are active entry only occurs when both active firms have an $\omega \leq 4$. I.e. 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 then when we do not. Indeed if, in a situation which 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 which does not allow for collusion, we would always expect more entry when collusion is allowed. This is a fact which has been largely been ignored in the literature and has important implications for the welfare analysis of collusive behavior.

**The Value Function and Investment.**

Figures 3a and 3b are plots of sections of the value function of the firm. Figure 3a assumes two firms are active; the firm we are studying (which has an $\omega$ which increases as we move away from the reader) and its competitor (whose $\omega$ increases as we move to the right on the graph). Figure 3b assumes three firms are active; the two whose $\omega's$ are plotted on the axis and a third firm at $\omega = 3$. Figures 4a provides the probability that our firm's research is successful when only two firm's are active (recall that this is a monotone transform of the firm's investment expenditures), while figure 4b 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 two figures provide the states at which a third 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 either when the incumbents have very low or very high values of $\omega$, but not when their $\omega's$ 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 $\omega$ is not monotone in its competitors state. The firm clearly prefers to be the only firm with a high $\omega$ for then it can force its competitor to exit, and charge monopoly prices until another entrant appears. However 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 competitors 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 the quality of its product are shared among the incumbents, and sometimes with 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 3a and 3b. In figure 3a (which, recall, assumes that only two firms active), there is only a narrow band of states in which the two 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 3b) the band in which collusion can be sustained but entry is not induced becomes quite a bit larger. Moreover, if the third firm's $\omega$ increases to 4 collusion is maintained and entry disappears altogether (see figure 4b).

24

As noted figure 4 is a transform of the *derivatives* of the value function in figure 3, and hence provides more details on the shape of that function. It is clear that the firm invests heavily when successful research has a large impact on the likelihood of the firm being in the collusive region, and then falls markedly. Moreover when the $\omega's$ 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. Note that 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 $\omega$ 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 $\omega's$ 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 and 4. 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 $\omega$ approaches it from below, for it is precisely when two firms have similar $\omega's$ that competition is most intense. In contrast when there was collusion the leading firm decreased its investment when the second firm approached it from below for if the two firms had similar $\omega's$ they could sustain collusion. Moreover if, in the noncollusive case, we move out along a diagonal with both firms having similar $\omega's$ 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.

25

## 4.2 The Collusive Industry; Descriptive and Welfare Analysis.

We used these policies to simulate 100,000 periods of industry evolution [27], 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.

We begin by describing the Nash equilibrium when there is no collusion (possibly because of an effective antitrust authority). The non-collusive industry is largely a "natural monopoly"; i.e. in 90 of the periods there is only one firm producing. The industry is also quite "stable"; the average length of a run with the same firm monopolizing production is 68 periods (the fourth panel in Table 1) .

The quality of the product of the monopolist 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 investment's 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 years and the median is three). This competitive phase typically ends with one firm pulling away from the others, and the others finding it unprofitable to compete with the superior product that firm developed and exiting the industry.

When we "take away the antitrust authority" and 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 in 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 $\omega$ (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

---

[27]Recall that a period here could be a fraction of a year, though in that case one would want to set the discount rate accordingly.

Table 1: *Market Structure.*

|  | (1) | (3) |
|---|---|---|
|  | MP<br>with<br>Collusion | MP<br>No<br>Collusion |
| **Percentage of equilibria with $n$ firms active** | | |
| $n=$ | | |
| 1 | 47.7 | 89.1 |
| 2 | 30.7 | 10.0 |
| (of these, with I=1) | (44.8) | n.r. |
| 3 | 21.1 | 0.8 |
| (of these, with I=1) | (77.3) | n.r. |
| 4 | 00.52 | 0.1 |
| (of these, with I=1) | (13.2) | n.r. |
| $\bar{n}$ | 1.74 | 1.12 |
| with entry | 5.4 | 2.9 |
| with exit | 5.3 | 2.2 |
| **Statistics On Lifespan Distribution** | | |
| Median | 3.0 | 3.0 |
| Mean | 32.7 | 51.1 |
| **Average Length of Runs for Given $(n, I)$ Couples.** * | | |
| $(n, I) =$ | | |
| (1,-) | 31.1 | 68.0 |
| (2,-) | 7.3 | 5.7 |
| (2,1) | 3.3 | n.r. |
| (2,0) | 3.3 | n.r. |
| (3,-) | 6.7 | 1.3 |
| (3,0) | 1.9 | n.r. |
| (3,1) | 5.2 | n.r. |
| (-,1) | 7.3 | n.r. |

* The $(2, -)$ row represents runs with the same two firms active, $(2, 1)$ runs with the same two firms colluding, etc. The $(-, 1)$ column represents a run of continuous collusion (though not necessarily between the same firms).

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 $\omega$ 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 $\omega$ of the initial incumbent continues to fall we will see two successive entry periods, and competition among three competitors with low $\omega's$. This type of competition, however, does not last long (on average less than 2 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 figure 3). 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 then 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 which induce entry. It is for this reason that length of the run with the same two firms colluding is typically quite short

28

(on average 3.3 periods). As noted when the two colluding firms states are high enough to induce entry and the new entrant invests successfully the collusive agreement 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 figure 4; apparently the Nash bargaining solution implies that too much of the increase in benefits obtained from higher $\omega$ would have to be shared with each other and the new entrant to make the investment itself worthwhile). It is important to realize, however, that when there is collusion we do observe three firms active with two of them producing high quality; when there is no collusion the only times we see three active firms is when all their products are low quality.

It is clear that an industry with collusive possibilities generates quite different evolutionary patterns than an industry without. The industry without collusive possibilities looks a lot like a natural monopoly which offers one relatively low quality product. Occasionally it is challenged by a new product and a competition ensues. The competition only continues, however, so long as none of the competing firms have 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 equilibria that we observe have tuples of states which are quite high. Indeed we frequently move to a tuple which is high enough to induce yet a third entrant. Consequently, not only are there more multi-product 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 then when there are none.

With these parameter values, then, allowing for collusion implies that

29

consumers will tend to have both more and higher quality products to choose from when collusion is allowed. Of course this need not make the consumers better off; it all depends on the prices that the 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 fewer and lower quality 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 five, so one can read the markups directly off this table. Consumer surplus differs with both this markup, and with $\omega$, so we require a separate calculation for that.

The most striking fact from table 2 is the difference between prices in collusive and non-collusive periods; conditional on the number of firms active that difference is *over fifty* percent. Moreover virtually all of 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 $\omega's$ that are quite high (see figure 2). The single firm wants to deter entry, but it can do this at an $\omega$ that is typically lower than the $\omega$'s needed to sustain collusion. That is the states of colluding firms are typically higher than the states of a monopolist; and this explains why prices are typically higher when there are colluding firms then 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 non-collusive industry are slightly higher. This is because there are so many more monopoly periods when collusion is not allowed. However, 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 $\omega$ than a monopolist in an industry which does not allow collusion. When we look at periods when two

Table 2: Prices and Benefits.

| | (1) MP with Collusion | (3) MP No Collusion |
|---|---|---|
| Prices* with $n$ firms active | | |
| $n=$ | | |
| 1 | 11.0 | 10.9 |
| 2 | 09.3 | 07.3 |
| of these, if $I = 1$ | 11.5 | n.r. |
| of these, if $I = 0$ | 07.5 | n.r. |
| 3 | 10.4 | 06.4 |
| of these, if $I = 1$ | 11.5 | n.r. |
| of these, if $I = 0$ | 6.6 | n.r. |
| average | 10.2 | 10.3 |
| Average Percentage Price Change When | | |
| Collusion begins | 47.0 | n.r. |
| Collusion ends | -48.0 | n.r. |
| Consumer and Producer Surplus** | | |
| Consumer Surplus | | |
| Mean | 22.8 | 19.8 |
| Std.Dev. | 06.3 | 07.2 |
| Producer Surplus | | |
| Mean | 34.5 | 38.1 |
| Std. Dev. | 11.1 | 10.7 |

*The Prices are the average, over periods, of the sales weighted average price.

** The consumer (producer) surplus measures are the mean and the standard deviation of the discounted sum of consumer (producer) benefits over a one hundred year period averaged over 1000 runs started at random draws from the ergodic distribution of states.

goods are offered one of two different scenarios unfold. 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 $\omega$ 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 panel of the table). Producer surplus is the discounted sum of; total profits minus total investment and entry costs plus any exit values, over a one-hundred year period. Consumer surplus is the discounted sum of consumer utility over the same period. The table contains the means and the standard deviations of the figures over a thousand separate samples from randomly drawn initial conditions [28].

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. However 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 (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).

---

[28]We ran one long run of length 100,000, and then broke the output up into a thousand subsamples of one hundred periods each.

32

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 which 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 would have worked with a more realistic distribution of utility functions, in particular a distribution which 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. The 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 which is preferred so much 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 a normal rate of return. Profits at any given state are higher when collusion is allowed, and apparently the inducement to enter generated by the higher profit potential has the effect of making the rents accruing to the sunk costs of the successful incumbents smaller.

Some final points. We have also calculated the equilibrium for the perfect cartel (a multi-product 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

33

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, which are very similar to those generated by the non-collusive 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. This does generate more surplus than in the collusive regime, and all the planner's surplus is allocated to consumers. In spite of this, Table 2 indicates that if a consumer orientated planner were given a choice between the non-collusive Nash solution (which from the point of view of concentration measures looks very close to the planner's solution), or the collusive Nash solution (which often 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 which has high enough sunk costs relative to demand for their to be only a small number of firms. On the other hand we have traditionally worried about collusion much more in industries with small numbers of active firms. Often the firms being investigated couch their defense in terms of the implications of "destructive competition". If the argument here is that if a virulent authority were to force "competitive" pricing, the industry would produce fewer (and possibly lower quality) products, and that this would hurt consumers, our calculations indicate that the firms could well be right.

# 5   Extensions

The paper has two goals. One was to present a particular model of collusive behavior and analyze its implications. The other was to provide a framework in which we could analyze the implications of collusion in a variety of dynamic settings. The framework has two important elements of dynamic interactions; (i)the strategic aspect which enables firms to condition their actions on the history of the interaction, and (ii) the structural aspect which implies the existence of state variables which affect the profit function and evolve over time in a manner that is partially controlled by the firms' actions. As noted in the paper, if one is only concerned with these two aspects of strategic interactions, our framework is quite flexible in the sense that we can change

34

the more detailed assumptions to better suit a given institutional setting, and then recompute and analyze the optimal policies.

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 at all the state variables of their competitors while during price wars states are perfectly observable. Assume now that at some particular period the firms agree on collusive prices (or quantities). These prices determine the allocation of the collusive rents among the firms. The prices (or quantity quotas) hold as long as there are no firms that demand to change them. However, if one of the firms gets good realizations of its investment efforts it may demand some changes in the collusive agreement. Changes that would reflect their new superior position. The superior position (high state variable) is however a private information and is not observable by rival firms. The demand for such changes may be accompanied by a threat of a price war. Without any threat there is no reason for the competitor firms to consider a reallocation of the collusive rents. If each demand for a new division of the collusive rents is conceded, firms may falsely announce on successes and on favorite states. Thus, in equilibrium it must be that some of these demands would be rejected. Rejection may follow by carrying out the threat of a price war. During such price wars the true states are revealed and firms may establish a new collusive agreement. The length of these price wars may be set in such a way to induce firms not to demand changes of the collusive agreement too frequently. The price war phase serves, in such a setup, also as a cost or a punishment for frequent demand for changes. While such a setup is beyond the scope of this paper, one can see that it is capable of generating price wars as part of the equilibrium behavior. The price wars reflects disagreement on the allocation of the collusive profits and as such the model and the resultant price wars resembles bargaining models with incomplete information about players characteristics that are capable to generate delays and strikes as part of the equilibrium behavior. It is also interesting to note that in a recent empirical study (Levenstein ,1997) of the 1885-1914 collusive behavior among the US bromine producers, some of the severe price wars occurred as part of a bargaining process in which one of the producers demand a renegotiation of the collusive agreement.

# 6 Concluding Remark

Standard anti-trust analysis of collusive behavior 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. However 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 anti-trust authority that is interested in maximizing social welfare should be worried about 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 then deterring it, since incentives generated by the collusive possibilities might well avert a monopoly situation. More generally a blanket "per se" aversion to collusion seems to be clearly off the mark.

The paper suggests that there is a need to revise standard anti-trust teaching and policy vis a vis collusion. We have also put 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 non-trivial way to be appropriate for settings in which asymmetric information is an important component of how collusion works. It seems clear, at least to us, that to understand the implications of collusion in any particular setting, we are going to have to mix the knowledge from several different subfields of economics; theory, institutional studies, econometric parameter estimates, and numerical analysis.

# References

Abreu, D., D. Pearce, and E. Stacchetti (1986), "Optimal Cartel equilibria with Imperfect Monitoring" *Journal of Economic Theory* 39(June), 251-269.

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

Bagwell, K. and R.W. Staiger (1997), "Collusion Over the Business Cy-

cle", *The Rand Journal of Economics*, Vol 28, pp 82-106.

Caplin, A. and B. Nalebuff (1991), "Aggregation and Imperfect Competition: the Existence of Equilibrium," *Econometrica* V, 59 pp. 25-59.

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

Ericson R. and A. Pakes, (1995) "Markov-Perfect Industry Dynamics: A Framework for Empirical Work" *Review of Economic studies*, 62, 53-82.

Fershtman, C. and E. Muller (1986), "Capital Investments and Price Agreements in Semi-Collusive Markets", *The Rand Journal of Economics* , vol.17, pp. 214-226.

Fershtman C. and N. Gandal (1994), "Disadvantageous Semicollusion" *International Journal of Industrial Organization*, vol.12, pp. 141-154.

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

Friedman J.W. and J.F. Thisse (1993), "Partial Collusion Fosters Minimum Product Differentiation" *The Rand Journal of Economics*, Vol. 24 pp 631-645.

Fudenberg, D., D. Levine, and E. Maskin (1994), "The Folk Theorem with Imperfect Public Information" *Econometrica*, 62(5) pp. 997-1039.

Hatiwanger, J. and J.E. Harrington (1991), "The Impact of Cyclical Demand Movements on Collusive Behavior", *The Rand Journal of Economics*, Vol. 22, pp. 89-106.

Harrington, J. (1989), "Collusion among Asymmetric Firms: The Case of Different Discount Factors" *International Journal of Industrial Organization*, 7, 289-307.

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

Green, E.J. and Porter, R. (1984), "Noncooperative Collusion Under Imperfect Price Competition" *Econometrica*, Vol. 52, pp.87-100.

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

Maskin, E. and J. Tirole (1988), "A Theory of Dynamic Oligopoly: I and II", *Econometrica*, 56, pp.549-600.

Maskin, E. and J. Tirole (1995), "Markov Perfect Equilibrium, I: Observable Actions", *mimeo*, Harvard University.

Nash, J.F. (1950), "The Bargaining Problem," *Econometrica* 18, pp 155-162.

Pakes, A. and P. McGuire (1994), "Computing Markov Perfect Nash Equilibrium: Numerical Implication of a Dynamic Differentiated Product Model", *The Rand Journal of Economics*, 25, pp. 555-589.

Pakes, A. and G. Gowrisankaran, and P. McGuire (1995), "Implementing the Pakes-McGuire Algorithm for Computing Markov Perfect Equilibria", *mimeo*, Yale University.

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

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

Rotemberg, J.J. and Saloner, G. (1986), " A supergame-theoretic model of price wars during booms" *American Economic Review*, Vol.76, pp. 390-407.

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

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

Staiger, R.W. and F.A. Wolak (1992), "Collusive Pricing with Capacity Constraint in the Presence of Demand Uncertainty." *The Rand Journal of Economics*, Vol. 23, pp. 203-220.
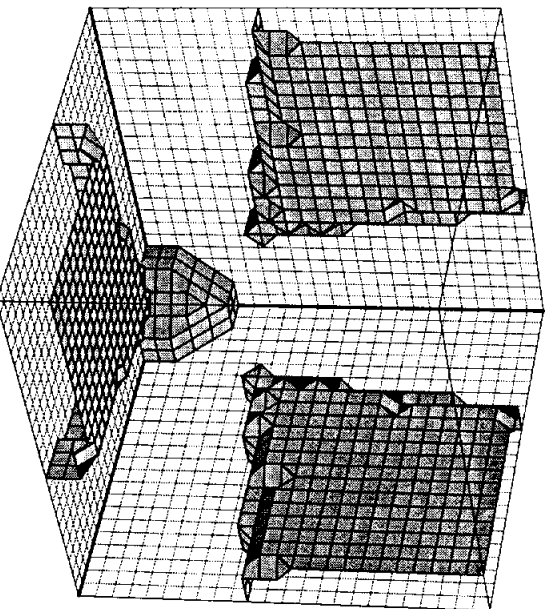
Starr, A.W. and Y.C. Ho (1969), "Nonzero-Sum Differential Games," *Journal of Optimization Theory and Application* Vol 3, pp.1984-208.

Stigler, G.J. (1964), "A Theory of Oligopoly" *Journal of Political Economy*, Vol. 72, pp.44-1.
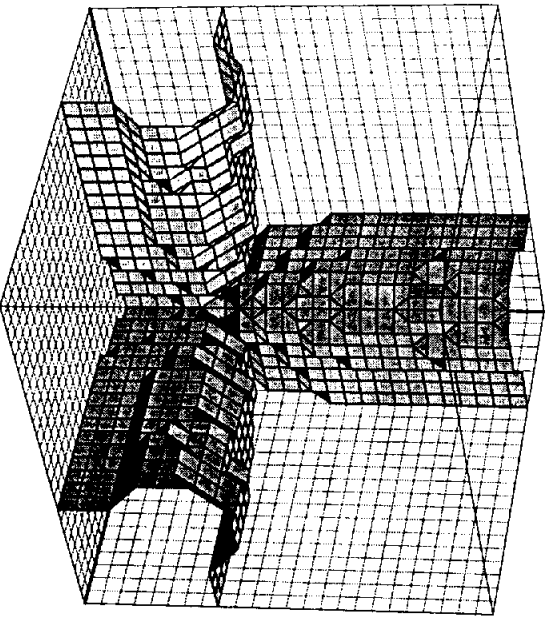
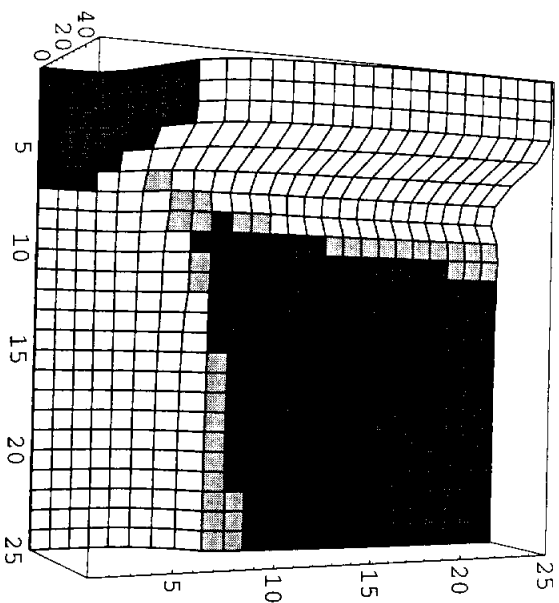Figure 1: Boundaries of Collusion Region



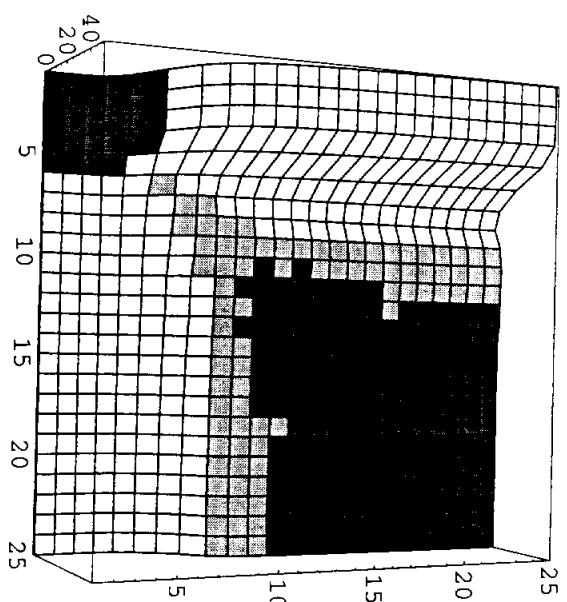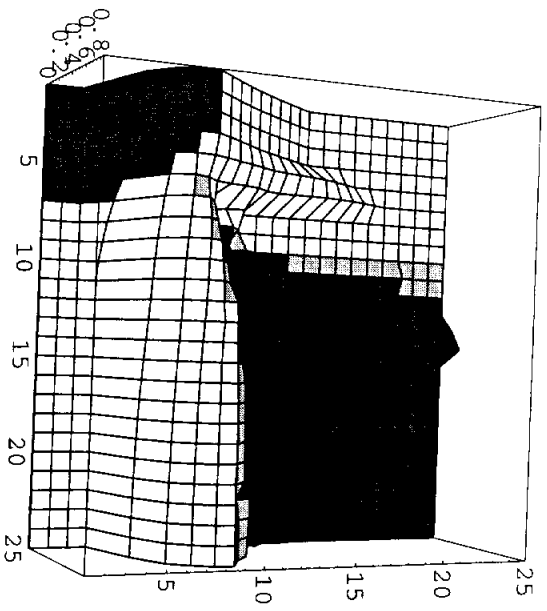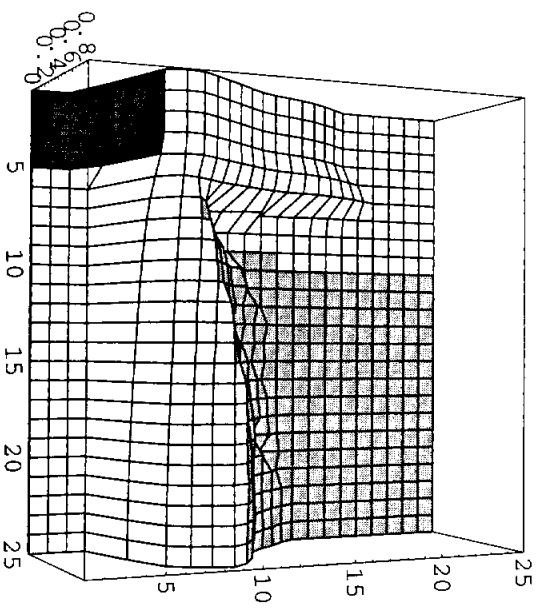Figure 2: Entry Region

Figure 3: Value Function, $\omega_3 = 0$



Figure 4: Value Function, $\omega_3 = 3$

**Figure 5: Investment**
**(Probability of Successful Research),  $\omega_3 = 3$**



**Figure 6: Investment**
**(Probability of Successful Research),  $\omega_3 = 4$**