

Game Theory and Industrial Organization

Kyle Bagwell and Asher Wolinsky*

April 3, 2000

Abstract

In this article, we consider how important developments in game theory have contributed to the theory of industrial organization. Our goal is not to survey the theory of industrial organization; rather, we consider the contribution of game theory through a careful discussion of a small number of topics within the industrial organization field. We also identify some points in which developments in the theory of industrial organization have contributed to game theory. The topics that we consider are: commitment in two-stage games and the associated theories of strategic-trade policy and entry deterrence; asymmetric-information games and the associated theories of limit pricing and predation; repeated games with public moves and the associated theory of collusion in markets with public demand fluctuations; mixed-strategy equilibria and purification theory and the associated theory of sales; and repeated games with imperfect monitoring and the associated theory of collusion and price wars. We conclude with a general assessment concerning the contribution of game theory to industrial organization.

1. Introduction

Game theory has become the standard language of industrial organization: the industrial organization theory literature is now presented almost exclusively in

*Bagwell: Columbia University (Economics Department and Graduate School of Business) and NBER; Wolinsky: Northwestern University (Economics Department). This article was written for the Handbook of Game Theory. We thank Rob Porter and Xavier Vives for helpful discussions. Financial support from the National Science Foundation is gratefully acknowledged.

terms of game theoretic models. But the relationship is not totally one-sided. First, the needs of industrial organization fed back and exerted a general influence on the agenda of game theory. Second, specific ideas that grew out of problems in industrial organization gained independent importance as game theoretic topics in their own right. Third, it is mostly through industrial organization that game theory was brought on large scale into economics and achieved its current standing as a fundamental branch of economic theory.

A systematic survey of the use of game theory in industrial organization would amount in fact to a survey of industrial organization theory. This is an enormous task that has been taken up by numerous textbooks.¹ The purpose of this article is not to survey this field, but rather to illustrate through the discussion of a small selection of subjects how some important developments in game theory have been incorporated into the theory of industrial organization and to pinpoint their contribution to this theory. We will also identify some points in which industrial organization theory raised a contribution to game theory. The models discussed are selected according to two criteria. First, they utilize a relatively major game theoretic idea. The second requirement is that the use of the game theoretic idea yields a relatively sharp economic insight.

Mathematical models in economics allow ideas to be expressed in a clear and precise way. In particular, they clarify the circumstances under which ideas are valid. They also facilitate the application of mathematical techniques, which sometimes yield insights that could not be obtained by simple introspection alone.² We will argue below that game theoretic models in industrial organization serve both of these functions.

As mentioned above, we do not intend to survey the field of industrial organization or the most important contributions to it. As a result, many important contributions and many influential contributors are not mentioned here. This should not be misinterpreted to suggest that these contributions are unimportant or that they are less important than those that were actually selected for the survey.

¹Tirole's (1988) comprehensive text is a standard reference.

²For example, some dynamic models begin with simple assumptions on, say, consumption and investment behavior that then give rise to a system that displays cyclical or even erratic aggregate behavior.

2. The role of commitment: an application of two-stage games

The role of commitment to future actions as a means of influencing rivals' behavior is a central theme in the analysis of oligopolistic competition. In a typical entry deterrence scenario, for example, an incumbent monopoly firm attempts to protect its market against entry of competitors by committing to post-entry behavior that would make entry unprofitable. In other scenarios, firms make partial commitments to future behavior through decisions on the adoption of technologies or through long-term contracts with their agents. The framework used in the literature for discussing these issues is that of a multi-stage game with subgame perfect equilibrium (Selten (1965)) as the solution concept.

THE BASIC MODEL

Two firms, 1 and 2, interact over two stages as follows. In the first stage the firms simultaneously choose the magnitudes, k_i , $i = 1, 2$. In the second, after observing the k_i 's, they choose simultaneously the magnitudes x_i , $i = 1, 2$. Firm i 's profit is given by the function $\pi_i(x_i, x_j; k_i)$, $i = 1, 2$, where $j \neq i$. A strategy for firm i , $[k_i, x_i(k_i, k_j)]$, prescribes a choice of k_i for stage 1 and a choice of x_i for stage 2, as a function of the k_i 's chosen in the first stage. A *subgame perfect equilibrium* (SPE) is a strategy pair, $[k_i^*, x_i^*(k_i, k_j)]$, $i = 1, 2$, such that: (A) for all (k_i, k_j) , $x_i^*(k_i, k_j) = \arg \max_x \pi_i[x, x_j^*(k_j, k_i); k_i]$; (B) $k_i^* = \arg \max_{k_i} \pi_i[x_i^*(k_i, k_j^*), x_j^*(k_j^*, k_i); k_i]$.

Thus, the x_i 's are the direct instruments of competition in that they enter the rival's profit directly, while the k_i 's have only an indirect effect. In many applications, the interpretation given is that the k_i 's represent productive capacities or technologies and the x_i 's describe quantities or prices of the final product. With the two-stage structure, k_i has a dual role: besides being a direct ingredient in the firm's profit, independently of the interaction, it also has a strategic role of influencing the rival's behavior in the second-stage subgame. The manner in which k_i affects x_j is credible in the sense of the SPE concept: k_i affects x_j only through shifting the second-stage equilibrium.

Perhaps the main qualitative result of this model, in its general form, is that the strategic role for k_i results in a distortion of its equilibrium level away from the level that would be optimal were x_j unaffected by k_i . When k_i is interpreted as capacity, this result means over or under investment in capacity as may be the case. The following proposition gives a precise statement of this result. Assume that

π_i , $i = 1, 2$, is differentiable, $\partial\pi_i/\partial x_j \neq 0$, there exists a unique SPE equilibrium, and x_i^* is a differentiable function of (k_i, k_j) .

Proposition 2.1: *If $\partial x_j^*(k_j^*, k_i^*)/\partial k_i \neq 0$, then*

$$\partial\pi_i[x_i^*(k_i^*, k_j^*), x_j^*(k_j^*, k_i^*); k_i^*]/\partial k_i \neq 0.$$

Proof. All the following derivatives are evaluated at the SPE point $[x_i^*(k_i^*, k_j^*), k_i^*]$, $i = 1, 2$. The first order condition for equilibrium condition (B) is $d\pi_i/dk_i = 0$, or

$$d\pi_i/dk_i \equiv (\partial\pi_i/\partial x_i)(\partial x_i^*/\partial k_i) + (\partial\pi_i/\partial x_j)(\partial x_j^*/\partial k_i) + (\partial\pi_i/\partial k_i) = 0.$$

Using the first order condition for condition (A), $\partial\pi_i/\partial x_i = 0$, we get

$$d\pi_i/dk_i = (\partial\pi_i/\partial x_j)(\partial x_j^*/\partial k_i) + (\partial\pi_i/\partial k_i) = 0$$

from which the proposition follows directly. ■

As a benchmark for comparison, consider the single-stage version of this game in which the firms choose (k_i, x_i) , $i = 1, 2$, simultaneously. The Nash equilibrium (Nash (1950)) of this game is (\bar{k}_i, \bar{x}_i) , $i = 1, 2$, such that $(\bar{k}_i, \bar{x}_i) = \arg \max_{k, x} \pi_i(x, \bar{x}_j; k)$.

Corollary: \bar{k}_i is in general different from k_i^* .

Proof. The first order condition for the equilibrium in the one-stage game is

$$\partial\pi_i(\bar{x}_i, \bar{x}_j; \bar{k}_i)/\partial k_i = 0.$$

Since $x_i^*(\bar{k}_i, \bar{k}_j) = \bar{x}_i$, it follows that

$$\partial\pi_i[x_i^*(\bar{k}_i, \bar{k}_j), x_j^*(\bar{k}_j, \bar{k}_i); \bar{k}_i]/\partial k_i = \partial\pi_i(\bar{x}_i, \bar{x}_j; \bar{k}_i)/\partial k_i = 0.$$

Therefore, it has to be that $(\bar{k}_i, \bar{k}_j) \neq (k_i^*, k_j^*)$, or else Proposition 2.1 will be contradicted. ■

DISCUSSION

From a conceptual point of view the two-stage oligopoly model is of course straightforward, and the main result of this model is an obvious implication of the SPE concept. The two-stage model does, however, provide a useful framework for discussing the role of commitment in oligopolistic competition. First, it embodies a clear notion of credibility. Second, it thereby serves to identify the features

that facilitate effective commitment: durable decisions that become observable before the actual interaction takes place. Third, it has been applied to a variety of specific economic scenarios and yielded interesting economic insights. The previous literature recognized the potential importance of commitment, but it did not manage to organize and understand the central idea in the clear form that the above model does.

An application to strategic trade theory

To see the type of economic insight that this model generates, consider its application to the theory of international trade (Brander and Spencer (1985)). Two firms, 1 and 2, are based in two different countries and export their products to a third country (the rest of the world). The product is homogenous, production costs are 0 and the demand is given by $p = 1 - Q$, where Q is the total quantity. The interaction unfolds in two stages. In the first stage, the two governments simultaneously choose excise tax (subsidy) rates, t_i , to be levied on their home firms. In the second stage, the tax rules are observed and the firms play a Cournot duopoly game: they simultaneously choose outputs q_i and the price is determined by $p = 1 - q_1 - q_2$. The effective cost functions in the second stage are $c_i(q_i; t_i) = t_i q_i$. The objective of firm i is maximization of its after-tax profit,

$$F_i(q_i, q_j; t_i) = (1 - q_i - q_j)q_i - t_i q_i.$$

The objective of each government is maximization of its country's "true" profit: the sum of its firm's profit and the tax revenue,

$$G_i(q_i, q_j; t_i) = F_i(q_i, q_j; t_i) + t_i q_i.$$

Since the government cares only about the sum, it chooses to tax only if the tax plays a strategic role and manipulates the second-stage competition in favor of its firm.

This application may be analyzed using the two-stage model developed above, although strictly speaking this is a slightly different case. The difference is that the stage-one commitments are now made by different parties (the exporting governments) than those who interact in stage two (the firms). However, the analysis remains the same. (The function G_i and the variables q_i and t_i in this case correspond to π_i , x_i and k_i in the general model above.) Solving for the SPE of this two-stage game, we get that the governments subsidize their firms: $t_1 = t_2 = -1/5$. In comparison to the equilibrium in the absence of government intervention, outputs and firms' profits are higher but countries' profits are

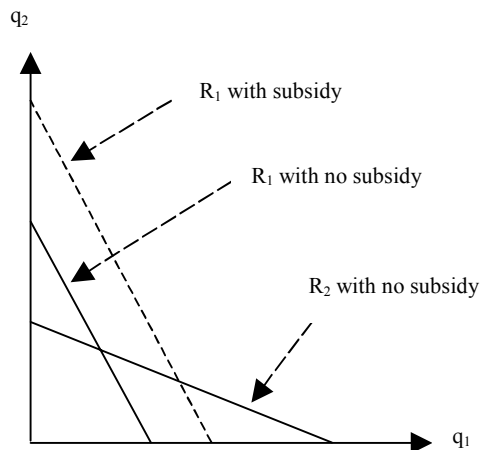


Figure 2.1:

lower. The intuition becomes more transparent by looking at the reaction functions, $R_i(q_j; t_i) = \arg \max_{q_i} F_i[q_i, q_j; t_i]$, depicted by Figure 2.1. The solid curves correspond to the case with no tax or subsidy. Their intersection point gives the second-stage equilibrium in this case. The dashed R_1 curve corresponds to a subsidy for firm 1, and its intersection with the R_2 curve gives the equilibrium when firm 1 is subsidized and firm 2 is not. The subsidy makes firm 1 more aggressive in the sense that, for any expectation that it might have regarding firm 2's output, it produces more than it would with no subsidy. This induces firm 2 to contract its output in equilibrium. Notice that, for a given output of firm 2, country 1's profit is higher with no subsidization, since the subsidy induces its firm to produce "too much." But the strategic effect on the other firm's output makes subsidization profitable and in equilibrium both governments offer export subsidies.

This is a striking insight that provides a clear and plausible explanation for export subsidies. To believe this explanation, one need not suppose that the governments see clearly through these strategic considerations. It is enough that they somehow think that subsidization improves their firms' competitive position. Also, despite its simplicity, this insight truly requires the game theoretic framework: it cannot be obtained without rigorous consideration of the strategic consequences of export policy.




Figure 2.2:

But further thought reveals that this insight is somewhat less convincing than it might have seemed at a first glance. Consider an alternative version of the model (Eaton and Grossman (1986)) in which the second-stage competition is a differentiated product Bertrand game: the firms simultaneously choose prices p_i and the demands are $q_i(p_i, p_j) = 1 - p_i + \alpha p_j$, with $0 < \alpha < 1$. Now,

$$F_i(p_i, p_j; t_i) = (1 - p_i + \alpha p_j)(p_i - t_i)$$

and

$$G_i(p_i, p_j; t_i) = F_i(p_i, p_j; t_i) + t_i(1 - p_i + \alpha p_j).$$

Repeating the above analysis for this case (now, the variable p_i corresponds to x_i in the general model), the reaction functions are $R_i(p_j; t_i) = \arg \max_{p_i} F_i[p_i, p_j; t_i]$. Figure 2.2 depicts the reaction functions in this case. Here, too, the dashed reaction function of firm 1 corresponds to a subsidy for firm 1. As before, the subsidy makes firm 1 more aggressive, inducing it to charge a lower price for any expectation it holds. But here this change in firm 1's position induces firm 2 to choose a lower price (as seen by comparing the two intersection points), and so an export subsidy now has a strategic cost, as it results in more aggressive behavior by the rival firm. Indeed, in the equilibrium of this scenario, the governments *tax*

their firms at level $t_1 = t_2 = \alpha^2(2 + \alpha)/[8 - 4\alpha^2 - \alpha^3]$, and prices and countries' profits are above their counterparts in the absence of intervention.

This, too, is a very clear insight. But, of course, the results here are almost the exact opposites of the results obtained above with the second-stage Cournot game. There are two views regarding the implications of this contrast. The more skeptical view maintains that the simple one-shot models of Cournot and Bertrand are only some sort of a parable. They are meant to capture the idea that in oligopolistic competition firms are aware that their rivals are also rational players who face similar decisions, and to point out that this sort of interaction might result in an inefficient outcome from the firms' perspective. But they are not meant to provide realistic descriptions of such competition. Thus, observations which depend on finer features of the structure of these models should not be regarded as true substantive insights. According to this view then, the only substantive insight here is that in principle there *might* be a strategic motive for the taxation/subsidization of exports. But the ambiguity of the results does not allow a useful prediction; in fact, it makes it hard to believe that this is a significant consideration in such scenarios.

A less skeptical view maintains that there is indeed a meaningful distinction between the sort of situations that are captured by the Cournot and Bertrand models. It can be argued that oligopolistic competition involves investments in production technologies (or capacities) followed by pricing decisions. The important strategic features of the Cournot model can be associated with the investment decisions, whereas the Bertrand model captures the pricing decisions.³ The results in the Cournot case thus rationalize strategic subsidization of research and development or other investment activities aimed at reducing cost or expanding capacity in export industries. So this view attributes to this analysis further content than the general insight that export subsidization might have a strategic role. It interprets the diverse predictions of the models of Cournot and Bertrand as reflecting important differences in the environment (e.g., regarding the age of the industry), which may understandably affect the outcome.

As mentioned above, the two-stage model reviewed here has been similarly employed in a number of different economic applications (e.g., capital investment, managerial incentive schemes). The sharp distinction between the predictions of the Cournot and Bertrand models appears in many of these applications.⁴ In the

³This understanding is related to the analysis of Kreps and Scheinkman (1983) and subsequent work.

⁴Bulow, Geanakoplos and Klemperer (1985) develop a general framework to which most

light of this, it is useful to emphasize that the qualitative effects described in the next application arise independently of the form of oligopoly competition.

Application to the question of entry deterrence

We close this section by reviewing briefly the development of the theoretical literature on entry deterrence that led to the adoption of the two-stage game framework discussed above. The manner in which this literature struggled with the concept of commitment may help to illustrate the nontrivial contribution of the above described framework toward improving the quality of this discussion. We do not expand here on the economic motivations of this literature, since these issues will be discussed in the next section.

Although earlier contributions to this literature took a variety of forms, it is convenient to present the ideas in the context of the two-stage framework of this section.⁵ An incumbent monopoly and a potential entrant interact over two periods. In the first (pre-entry) period, the incumbent selects a price, which is observed by the entrant. In the second (post-entry) period, the entrant decides whether or not to enter. Entry entails a fixed cost, and the incumbent's profit in the post-entry interaction with the entrant is lower than its profit as a continuing monopoly.

The earlier literature explored a particular model of this form and developed the notion of a *limit price*. In the context of this model, the incumbent limit prices when it chooses a relatively low price, typically lower than the regular monopoly price, that would render entry unprofitable, if this price were to prevail as well in the post-entry period. The potential entrant then responds by staying out. This model, however, entails the seemingly implausible assumption that the incumbent would choose to maintain its pre-entry price in the event that entry actually occurred. Furthermore, if we are unwilling to make this assumption, then it is no longer clear why the pre-entry price should affect the expected profit from entry and thus the entry decision itself. If we think, for example, that the post-entry interaction in fact takes the form of a standard duopoly game, and that the post-entry demand and cost functions are independent of the pre-entry price, then the potential entrant's expected duopoly profit should also be independent of the pre-entry price. It is therefore doubtful that the entry of rational competitors

of these applications belong. They also coined the terms "strategic substitutes" and "strategic complements" to describe the cases of downward- and upward-sloping reaction functions, respectively. See also Fudenberg and Tirole (1984).

⁵Important contributions to the early literature on limit pricing include those by Bain (1949) and Modigliani (1958).

can be blocked in this manner. This suggests that limit pricing emerges as part of a credible entry deterrence strategy, only if some alternative mechanism (other than price commitment) is identified that links the incumbent's pre-entry price to the potential entrant's expected profit from entry.

Motivated by this understanding, the next step in the development of this theory (Spence (1977)) introduced the possibility that the incumbent selects a level of capacity in the pre-entry period. An investment in capacity is plausibly irreversible, and so an investment of this kind is a natural means through which an incumbent may credibly commit to be an active participant in the market. In particular, the idea was that entry would be deterred when an incumbent invested significantly in capacity, if the incumbent were to threaten that it would utilize its capacity to the fullest extent in the post-entry interaction. The notion of a limit price gets here a different meaning. The pre-entry price is no longer a strategic instrument for blocking entry. But if the entry deterring investment level reduces the incumbent's marginal cost relative to that of an unthreatened monopoly, the entry threat might have the effect of lowering the incumbent's pre-entry price. This would have the appearance of a limit price, but it is actually only a by-product of the entry deterring investment.

While Spence's model identified capacity as a plausible pre-entry commitment variable, it still did not evaluate the credibility of the threatened utilization of the installed capacity. In fact, the threatened entry deterring output is not always credible in the SPE sense—the equilibrium in a post-entry duopoly game does not necessarily entail utilization of the capacity installed as a threat. This shortcoming was addressed by the next step of this theory (Dixit (1980)) which introduced a formal two-stage game model, of the family discussed in this section, with SPE as the solution concept. The incumbent's threat backed by its pre-entry investment is credible in the sense that it manifests itself through its effect on the equilibrium of the post-entry duopoly game. Appropriate versions of this model thus explain excessive investment in capacity and the associated low pre-entry price as credible responses to entry threats.

From the viewpoint of pure game theory, the final model that capped this literature is rather straightforward. But the extent of its contribution, even if only to sharpen the relationship between price, capacity and deterrence, should be evident from looking at the long process of insightful research that led to that point.

3. Entry deterrence and predation: applications of sequential equilibrium

It is widely agreed that unhindered exercise of monopoly power generally results in inefficient resource allocation and that anti-trust policy aimed at the prevention of monopolization is therefore a legitimate form of government intervention in the operation of markets. A major concern of anti-trust policy has been the identification and prevention of practices that lead to monopolization of industries. This concern has motivated a large body of theory aimed at understanding and classifying the different forms that monopolization efforts might take and their economic consequences.

Monopolization takes a variety of forms ranging from more cooperative endeavors, like merger and cartelization, to more hostile practices, like entry deterrence and predation. The last part of the previous section described some of the developments in the understanding of entry deterrence by means of pricing and preemptive investment. The related notion of predation refers to attempts to induce the exit of competitors by using similar aggressive tactics. The treatment of predation and entry deterrence raises some subtle issues, both in theory and in practice, since it is naturally difficult to distinguish between “legitimate” competitive behavior that enhances efficiency and anti-competitive behavior that ultimately reduces welfare.

In fact, it has been often argued that predation or entry deterrence through aggressive pricing behavior is not a viable strategy among rational firms. The implication is that instances of aggressive price cutting should not be interpreted along these lines. The logic of this argument was explained in the previous section: a credible predatory or entry deterring activity must create a meaningful *commitment link* to the behavior of this firm in its future interactions. But, with rational players, an aggressive pricing policy does not in itself plausibly constitute such a commitment. This argument leads to the following conclusion: when aggressive pricing appears in instances of entry deterrence (or predation), it is a by-product of a strategic investment in capacity, rather than a strategic instrument in its own right.

The following discussion exposes a significant limitation to this conclusion. It shows that, when informational differences about cost (or other parameters that affect profit) are important, it is possible to revive the traditional view of the limit price literature that pricing policies may serve as direct instruments of monopolization. The idea is that, in the presence of such asymmetric information,

prices might also transmit information and as such play a direct role in influencing the entry or exit decisions of rivals.

LIMIT PRICING UNDER ASYMMETRIC INFORMATION

Milgrom and Roberts (1982a) consider the classic scenario of a monopoly incumbent facing a single potential entrant. The novel feature of their analysis is that the incumbent has some private information regarding its costs of production. This information is valuable to the entrant, as its post-entry profit is affected by the incumbent's level of cost: the lower is the incumbent's cost, the lower is the entrant's profit from entry. Thus, in place of the commitment link studied by the earlier literature, Milgrom and Roberts propose an *informational link* between the incumbent's pre-entry behavior and the entrant's expected post-entry profit.

The situation can be modeled as a signaling game: the entrant attempts to infer the cost information by observing the incumbent's pre-entry pricing, while the incumbent chooses its price with the understanding that this choice may affect the prospect of entry. The incumbent would like the entrant to think that its costs are low to make entry seem less profitable. As in other signaling models, the equilibrium price-signal therefore may be distorted away from its myopic level. In the present context, the price is distorted if it differs from the myopic monopoly price (i.e., the price that would prevail in the absence of the entry threat). To correspond to the original limit price conjecture, the equilibrium price has to be distorted downwards from the monopoly level. But since lower costs naturally lead to lower monopoly prices, the downwards distortion is indeed the expected result in a signaling scenario.

The details of the game are as follows.⁶ The players are an incumbent monopoly firm and a potential entrant. The interaction spans two periods: the pre-entry period in which the entrant observes the incumbent's behavior and contemplates the entry decision, and the post-entry period after the entry decision is resolved in either way. The incumbent is one of two possible types of firm differing in their per-unit cost of production: type $t \in \{L, H\}$ has unit cost $c(t)$, where $c(L) < c(H)$ so that L and H stand for "low" and "high," respectively. In the pre-entry period, only the incumbent knows its type. It chooses a pre-entry price, p , which the entrant observes and on the basis of which decides whether to enter. The incumbent's profit in the pre-entry period is $\Pi(p, t) = [p - c(t)]D(p)$, where D is a well-behaved downwards-sloping demand function. We abstract from the details

⁶The discussion here expands on previous presentations by Bagwell and Ramey (1988) and Fudenberg and Tirole (1986).

associated with the play of firms in the post-entry period, and simply summarize the outcomes of that period. If the entrant does not enter, then its profit is 0 and the incumbent remains a monopoly and earns $\pi^m(t)$; if the entrant does enter, it learns the incumbent's cost and the resulting post-entry duopoly profits are $\pi^d(t)$ for the incumbent and $\pi^e(t)$ to the entrant. It is assumed that entry reduces the incumbent's post-entry profit, $\pi^m(t) > \pi^d(t)$, and that the entrant can recover any fixed costs associated with entry only against the high-cost incumbent, $\pi^e(H) > 0 > \pi^e(L)$. Note that $\pi^m(t)$ admits a variety of interpretations: it might be simply the discounted maximized value of $\Pi(p, t)$, or it might pertain to a different length of time and/or reflect some further future interactions.

The game theoretic model is then a simple sequential game of incomplete information. The formal description is as follows. Nature chooses the incumbent's type $t \in \{L, H\}$ with probability b_t^0 , where $b_L^0 + b_H^0 = 1$. The incumbent's strategy is a pricing function $P : \{L, H\} \rightarrow [0, \infty)$. The entrant's belief function $b_t : [0, \infty) \rightarrow [0, 1]$ describes the probability it assigns to type t , given the incumbent's price p . Of course, for all p , $b_L(p) + b_H(p) = 1$. A strategy for the entrant is a function $E : [0, \infty) \rightarrow \{0, 1\}$ that describes the entry decision as a function of the incumbent's price, where "1" and "0" represent the entry and no-entry decisions, respectively. The payoffs as functions of price $p \in [0, \infty)$, entry decision $e \in \{0, 1\}$, and cost type $t \in \{L, H\}$ are: for the incumbent,

$$V(p, e, t) = \Pi(p, t) + e\pi^d(t) + (1 - e)\pi^m(t)$$

and for the entrant,

$$u(p, e, t) = e\pi^e(t).$$

It is convenient to introduce special notation for the entrant's expected payoff evaluated with its beliefs. Letting b denote the pair of functions (b_L, b_H) ,

$$U(p, e, b) = b_L(p)u(p, e, L) + b_H(p)u(p, e, H).$$

Observe that $U(p, 0, b) = 0$ and $U(p, 1, b) = b_L(p)\pi^e(L) + b_H(p)\pi^e(H)$.

The solution concept is *sequential equilibrium* (Kreps and Wilson (1982a)) augmented by the "intuitive criterion" refinement (Cho and Kreps (1987)). For the present game, a sequential equilibrium is a specification of strategies and beliefs, $\{P, E, b\}$, satisfying three requirements:

(E1) Rationality for the incumbent:

$$P(t) \in \arg \max_p V(p, E(p), t), \quad t = L, H$$

(E2) Rationality for the entrant:

$$E(p) \in \arg \max_e U(p, e, b(p)) \quad \text{for all } p \geq 0$$

(E3) Bayes-Consistency:

$$\begin{aligned} P(L) = P(H) &\text{ implies } b_L(P(L)) = b_L^0. \\ P(L) \neq P(H) &\text{ implies } b_L(P(L)) = 1, \quad b_L(P(H)) = 0 \end{aligned}$$

As (E3) indicates, there are two types of sequential equilibria. In a *pooling equilibrium* ($P(L) = P(H)$), the entrant learns nothing from the observation of the equilibrium price, and so the posterior and the prior beliefs agree; whereas in a *separating equilibrium* ($P(L) \neq P(H)$), the entrant is able to infer the incumbent's cost type upon observing the equilibrium price.

For this game, sequential equilibrium places no restrictions on the beliefs that the entrant holds when a deviant price $p \notin \{P(L), P(H)\}$ is observed. For example, the analyst may specify that the entrant is very optimistic and infers high costs upon observing a deviant price. In this event, the incumbent may be especially reluctant to deviate from a proposed equilibrium, and so it becomes possible to construct a great many sequential equilibria. The set of sequential equilibria here will be refined by imposing the following “plausibility” restriction on the entrant's beliefs following a deviant price:

(E4) Intuitive beliefs:

$$\begin{aligned} b_t(p) &= 1 \text{ if } \text{for } t \neq t' \in \{L, H\}, \\ V(p, 0, t) &\geq V[P(t), E(P(t)), t] \text{ and } V(p, 0, t') < V[P(t'), E(P(t')), t'] \end{aligned}$$

The idea is that an incumbent of given type would never charge a price p that, even when followed by the most favorable response of the entrant, would be less profitable than following the equilibrium. Thus, if a deviant price is observed that could possibly improve upon the equilibrium profit only for a low-cost incumbent, then the entrant should believe that the incumbent has low costs. In what follows, we say that a triplet $\{P, E, b\}$ forms an *intuitive equilibrium* if it satisfies (E1)-(E3) and (E4).

Before turning to the equilibrium analysis, we impose some more structure on the payoffs. The first assumption is just a standard technical one, but the second injects a further meaning to the distinction between the types by ensuring that the low-cost type is also more eager to prevent entry:

(A1) The function Π is well behaved:

$\exists \tilde{p} > c(H)$ such that $D(p) > 0$ iff $p < \tilde{p}$, and Π is strictly concave and differentiable on $(0, \tilde{p})$.

(A2) The low-cost incumbent loses more from entry:

$$\pi^m(L) - \pi^d(L) \geq \pi^m(H) - \pi^d(H).$$

Assumption (A2) is not obviously compelling. It is natural to assume that the low-cost incumbent fares better than the high-cost one in any case, $\pi^m(L) > \pi^m(H)$ and $\pi^d(L) > \pi^d(H)$, but this does not imply the assumed relationship. This assumption is satisfied in a number of well-behaved specifications of the post-entry duopolistic interaction, but it might be violated in other standard examples.

We observe next that a low-cost incumbent is more attracted to a low pre-entry price than is a high-cost incumbent, since the consequent increase in demand is less costly for the low-cost incumbent. Formally, for $p, p' \in (0, \tilde{p})$ such that $p < p'$,

$$\Pi(p, L) - \Pi(p, H) \equiv [c(H) - c(L)]D(p) > [c(H) - c(L)]D(p') \equiv \Pi(p', L) - \Pi(p', H).$$

This together with assumption (A2) immediately imply the following *single crossing property* (SCP):

For any $p < p'$ and $e \leq e'$, if $V(p, e, H) = V(p', e', H)$, then $V(p, e, L) > V(p', e', L)$.

Under (SCP), if a high-cost incumbent is indifferent between two price-entry pairs, then a low-cost incumbent prefers the pair with a lower price and a (weakly) lower rate of entry. In particular, to deter entry the low-cost incumbent would be willing to accept a deeper price cut than would the high-cost incumbent. As is true throughout the literature on signaling games, characterization and interpretation of the equilibria is straightforward when the preferences of the informed player, here the incumbent, satisfy an appropriate version of (SCP).

Let $p_t^m = \arg \max_p \Pi(p, t)$. This is the (myopic) monopoly price of an incumbent of type t . Under our assumptions, it is easily confirmed that the low-cost

monopoly price is less than that of the high-cost incumbent: $p_L^m < p_H^m$. Consider now the set of prices p such that

$$V(p, 0, H) \leq V(p_H^m, 1, H).$$

The concavity of Π in p assures that this inequality holds outside an interval $(\underline{p}, \bar{p}) \subset [0, \hat{p}]$, and its reverse holds inside the interval. Thus, \underline{p} and \bar{p} are the prices which give the high-cost incumbent the same payoff when it deters entry as when it selects the high-cost monopoly price and faces entry. Since entry deterrence is valuable, it follows directly that $\underline{p} < p_H^m < \bar{p}$.

Let \tilde{b}_L denote the belief that would make the entrant exactly indifferent with respect to entry. It is defined by

$$\tilde{b}_L \pi^e(L) + (1 - \tilde{b}_L) \pi^e(H) = 0.$$

Proposition 3.1: (i) *There exists a separating intuitive equilibrium.*

(ii) *If $p_L^m \geq \underline{p}$, then any separating intuitive equilibrium, $\{P, E, b\}$, satisfies $\underline{p} = P(L) < P(H) = p_H^m$ and $E(P(L)) = 0 < 1 = E(P(H))$.*

If $p_L^m < \underline{p}$, then any separating intuitive equilibrium, $\{P, E, b\}$, satisfies $p_L^m = P(L) < P(H) = p_H^m$ and $E(P(L)) = 0 < 1 = E(P(H))$.

(iii) *If $p_L^m \geq \underline{p}$, and $b_L^0 \geq \tilde{b}_L$, then for every $p \in [\underline{p}, p_L^m]$, there exists an intuitive pooling equilibrium in which $P(L) = P(H) = p$.*

(iv) *In any intuitive pooling equilibrium, $P(L) = P(H) \in [\underline{p}, p_L^m]$ and $E(P(L)) = 0$.*

The proof is relegated to the appendix. The proposition establishes uniqueness of the separating equilibrium outcome. Since the high-cost incumbent faces entry in the separating equilibrium, its equilibrium price must coincide with its monopoly price p_H^m . Otherwise, it would clearly profit by deviating to p_H^m from any other price at which it anyway faces entry. The case $p_L^m \geq \underline{p}$ corresponds to a relatively small cost differential between the two types of incumbent. It is the more interesting case, since the separating equilibrium price quoted by the low-cost incumbent is then distorted away from its monopoly price p_L^m . The case $p_L^m < \underline{p}$ corresponds to relatively large cost differential which renders p_L^m a dominated choice for the high-cost incumbent and hence removes the tension associated with the high-cost incumbent's incentive to mimic the low-cost price.

When the prior probability of the low-cost type is sufficiently large, $b_L^0 \geq \tilde{b}_L$, there are also pooling equilibria. These equilibria do not exist when $b_L^0 < \tilde{b}_L$, since

at a putative pooling equilibrium the entrant would choose to enter and hence the incumbent would profit from deviating to its monopoly price p_t^m .

Both the separating and the pooling equilibria exhibit limit price behavior, but these patterns are qualitatively very different. The separating equilibrium exhibits limit pricing in the sense that, for certain parameter values, $P(L) < p_L^m$. But the separating equilibrium differs from the traditional limit price theory in the important sense that equilibrium limit pricing does not deter entry, which occurs under the same conditions (namely, when the incumbent has high costs) that would generate entry in a complete-information setting. The effect of the limit price on entry is through the cost information that it credibly reveals to the entrant. Limit pricing also occurs in the pooling equilibria. The high-cost incumbent now practices limit pricing, as $P(H) < p_H^m$ in any pooling equilibrium, and the low-cost incumbent also selects a limit price, as $P(L) < p_L^m$ in all these equilibria save the one in which pooling occurs at p_L^m . In contrast with the limit pricing of the separating equilibrium, and in accordance with the traditional notion, the limit price here does deter entry. The rate of entry is lower than would occur under complete information, since the high-cost incumbent is able to deter entry when it pools its price with that of the low-cost incumbent.

Earlier literature on the traditional notion of limit pricing associated with this practice a welfare trade-off: lower prices generate immediate welfare gains but deter or reduce entry and thus lead to future welfare losses. The form of limit pricing that arises under the separating equilibrium is actually beneficial for welfare, since the low-cost incumbent signals its information with a low price and this does not come at the expense of entry. Instead, it is the pooling equilibria that exhibit the welfare trade-off that the earlier literature associated with limit pricing. While the low pre-entry prices tend to improve welfare, the reduction in entry lowers welfare in the post-entry period, as compared to the welfare that would be achieved in a complete-information setting.

The set of equilibria may be further refined with a requirement that the selected equilibrium is Pareto efficient for the low- and high-cost incumbent among the set of intuitive equilibria. When pooling equilibria exist (the conditions of part (iii) of the proposition hold), then the pooling equilibrium in which the low-cost monopoly price is selected is the efficient one for the low- and high-cost incumbent in the relevant set. This equilibrium gives the low-cost incumbent the maximum possible payoff. It also offers a higher payoff to the high-cost incumbent than occurs in the separating equilibrium, since $p_L^m \geq \underline{p}$ implies that $V(p_L^m, 0, H) \geq V(\underline{p}, 0, H) = V(p_H^m, 1, H)$.

PREDATION

Generally speaking, a firm practices predatory pricing if it charges “abnormally” low prices in an attempt to induce exit of its competitors. The ambiguity of this definition is not incidental. It reflects the inherent difficulty of drawing a clear distinction between legitimate price competition and pricing behavior that embodies predatory intent. Indeed, an important objective of the theoretical discussion of this subject is to come up with relatively simple criteria for distinguishing between legitimate price competition and predation.

Exit inducing behavior is of course closely related to entry deterring behavior and, indeed, a small variation on the limit pricing model presented above can also be used to discuss predation. Consider, then, the following variation on the limit price game. In the first period, both firms (referred to as the “predator” and the “prey”) are in the market and choose prices simultaneously. Then the prey, who must incur a fixed cost if it remains in the market, decides whether or not to exit. Finally, in the second period, the active firms again choose prices. If the prey exits, the predator earns monopoly profit; on the other hand, if the prey remains in the market, the two firms earn some duopoly profits. In equilibrium, the prey exits when its expected period-two profit is insufficient to cover its fixed costs. Clearly, in any SPE of the complete-information version of this game, no predation takes place, as the prey’s expectation is independent of the predator’s first-stage price. However, when the prey is uncertain about the predator’s cost, as in the above limit pricing model, then an informational link appears between the predator’s first-period price and the prey’s expected profit from remaining in the market. Recognizing that the prey will base its exit decision upon its inference of the predator’s cost type, the predator may price low in order to signal that its costs are low and thus induce exit. The equilibria of this model are analogous to those described in Proposition 3.1, with exit occurring under the analogous circumstances to those under which entry was deterred. This variation provides an equilibrium foundation for the practice of predatory pricing, in which predation is identified with low prices that are selected with the intention of affecting the exit decision.

From a welfare standpoint, the predation that occurs as part of a separating equilibrium is actually beneficial. Predation brings the immediate welfare benefit of a lower price, and it induces the exit of a rival in exactly the same circumstances as would occur in a complete-information environment. When the game is expanded to include an initial entry stage (Roberts (1985)), however, a new wrinkle appears, as the rational anticipation of predatory signaling may deter the

entry of the prey, resulting in a possible welfare cost.

DISCUSSION

The notion that limit pricing can serve to deter entry has a long history in industrial organization, with a number of theoretical and empirical contributions.⁷ The signaling model of entry deterrence contributed to this literature a number of new theoretical insights. First and foremost, it identified two patterns of rational behavior that may be interpreted in terms of the “anti-competitive” practices of entry deterrence and predation. One pattern, exemplified by the pooling equilibria, exhibits anti-competitive behavior in its traditional meaning of eliminating competition that would otherwise exist. The other pattern, exemplified by the separating equilibrium, takes the appearance of anti-competitive behavior but does not exhibit the traditional welfare consequences of such behavior. These observations have a believable quality to them both because the main element of this model is asymmetric information, which is surely often present in such situations, and because the pooling and separating equilibria have natural and intuitive interpretations. Furthermore, continued research has shown that the basic ideas of this theory are robust on many fronts.⁸

Even if these insights had been on some level familiar prior to the introduction of this model, and this is doubtful, they surely had not been understood as implications of a closed and internally consistent argument. In fact, it is hard to envision how these insights could be derived or effectively presented without the game theoretic framework. So this part of the contribution, the generation and the crisp presentation of these insights, cannot be doubted.

Still there is the question of whether these elegant insights change significantly our understanding of actual monopolization efforts. Here, it is useful to distinguish between qualitative and quantitative contributions. Certainly, as the discussion above indicates, the signaling model of entry deterrence offers a significant qualitative argument that identifies a possible role for pricing in anti-competitive

⁷A recent empirical analysis is offered by Kadiyali (1996), who studies the U.S. photographic film industry and reports evidence that is consistent with the view that the incumbent (Kodak) selected a low price and a high level of advertising in the presence of a potential entrant (Fuji).

⁸There are, however, a couple of variations which alter the results in important ways. First, as Harrington (1986) shows, if the entrant’s costs are positively correlated with those of the incumbent, then separating equilibria entail an upward distortion in the high-cost incumbent’s price. Second, Bagwell and Ramey (1991) show that, when the industry hosts several incumbents who share private information concerning industry costs, a focal separating equilibrium exists that entails no pricing distortions whatsoever.

behavior. It is more difficult, however, to assess the extent to which this argument will lead to a quantitative improvement in our understanding of monopolization efforts. For example, is it possible to identify with confidence industries in which behavior is substantially affected by such considerations? Can we seriously hope to estimate (perhaps even roughly) the quantitative implication of this theory? We do not have straightforward answers to these questions. The difficulties in measurement would make it quite hard to distinguish between the predictions of this theory and others.

Even the qualitative features of the argument may be of important use in the formulation of public policy. For example, U.S. anti-trust policy aims at curbing monopolization and specifically prohibits predatory pricing. However, the exact meaning of this prohibition as well as the manner in which it is enforced have been subject to continued review over time by the government and the courts. Two ongoing debates that influence the thinking on this matter are as follows: Is predation a viable practice in rational interaction? If the possibility of predation is accepted, what is the appropriate practical definition of predation? The obvious policy implication in case predation is not deemed to be a viable practice is that government intervention is not needed. In the absence of a satisfactory framework that can supply precise answers, these policy decisions are shaped by weighing an array of incomplete arguments. Historically, some of the most influential arguments have been developed as simple applications of basic price theoretic models. Prominent among these are the “Chicago School” arguments that deny the viability of predation (McGee (1958)) and the Areeda-Turner Rule (1975) that associates the act of predation with a price that lies below marginal cost.

With this in mind, there is no doubt that the limit pricing model enriches the arsenal of arguments in a significant way. First, it provides a theoretical framework that clearly establishes the viability of predatory behavior among rational competitors. Second, it raises some questions regarding the practical merit of cost-based definitions of predation, like the Areeda-Turner standard: it shows that predation might occur under broader circumstances than such standards admit. In a world in which a government bureaucrat or a judge has to reach a decision on the basis of imprecise impressions, arguments that rely on the logic of this theory may well have important influence.⁹

⁹Other game theoretic treatments of predation, like the reputation theories of Kreps and Wilson (1982b) and Milgrom and Roberts (1982b) or the war of attrition perspective of Fudenberg and Tirole (1987), also provide similar intellectual underpinning for government intervention in curbing predation.

4. Collusion: an application of repeated games

One of the main insights drawn from the basic static oligopoly models concerns the inefficiency (from the firms' viewpoint) of oligopolistic competition: industry profit is not maximized in equilibrium.¹⁰ This inefficiency creates an obvious incentive for oligopolists to enter into a collusive agreement and thereby achieve a superior outcome that better exploits their monopoly power. Collusion, however, is difficult to sustain, since typically each of the colluding firms has incentive to opportunistically cheat on the agreement. The sustenance of collusion is further complicated by the fact that explicit collusion is often outlawed. In such cases collusive agreements can not be enforced with reference to legally binding contracts. Instead, collusive agreements then must be "self enforcing": they are sustained through an implicit understanding that "excessively" competitive behavior by any one firm will soon lead to similar behavior by other firms.

Collusion is an important subject in industrial organization. Its presence in oligopolistic markets tends to further distort the allocation of resources in the monopolistic direction. For this reason, public policy toward collusion is usually antagonistic. While there is both anecdotal and more systematic evidence on the existence of collusive behavior, the informal nature and often illegal status of collusion makes it difficult to evaluate its extent. But regardless of the economy-wide significance of collusion, this form of behavior is of course of great significance for certain markets.

The main framework currently used for modeling collusion is that of an infinitely repeated oligopoly game. Since the basic tension in the collusion scenario is dynamic—a colluding firm must balance the immediate gains from opportunistic behavior against the future consequences of its detection—the analysis of collusion requires a dynamic game which allows for history dependent behavior. The repeated game is perhaps the simplest model of this type.

The earlier literature, which preceded the introduction of the repeated game model, recognized the basic tension that confronts colluding firms and the factors that affect the stability of collusive agreements (see, e.g., Stigler (1964)). In particular, this literature contained the understanding that oligopolistic interaction sometimes results in collusion sustained by threats of future retaliation, and at other times results in non-collusive behavior of the type captured by the equilibria

¹⁰This result appears in an extreme form in the case of the pure Bertrand model, in which two producers of a homogenous product who incur constant per-unit costs choose their prices simultaneously. The unique equilibrium prices are equal to marginal cost and profits are zero.

of the static oligopoly models. But since the formal modeling of this phenomenon requires a dynamic strategic model, which was not then available in economics, this literature lacked a coherent formal model.¹¹

The main contribution of the repeated game model of collusion was the introduction of a coherent formal model. The introduction of this model offers two advantages. First, with such a model, it is possible to present and discuss the factors affecting collusion in a more compact and orderly manner. Second, the model enables exploration of more complex relations between the form and extent of collusive behavior and the underlying features of the market environment. A contribution of this type is illustrated below by a simple model that characterizes the behavior of collusive prices in markets with fluctuating demand (Rotemberg and Saloner (1986)). We have selected to feature this application, since it draws a clear economic insight by utilizing closely the particular structure of the repeated game model of collusion.

PRICE WARS DURING BOOMS.

Two firms play the Bertrand pricing game repeatedly in the following environment. In each period t the market demand is inelastic up to price 1 at quantity a^t ,

$$Q(P) = \begin{cases} 0 & \text{if } P > 1 \\ a^t & \text{if } P \leq 1 \end{cases}$$

The production costs are zero. The a^t 's are i.i.d. random variables which take the values H and L , where $H > L$ and $\text{Prob}\{a^t = H\} = w = 1 - \text{Prob}\{a^t = L\}$. Within each period, events unfold in the following order. First, a^t is realized and observed by the firms. Second, the firms choose prices, $p_i^t \in [0, 1]$, $i = 1, 2$, simultaneously. Third, the instantaneous profits, $\pi_i^t(p_i^t, p_j^t)$, are determined and distributed. The firm with the lower price gets the entire market, and when prices are equal the market is split equally. That is,

$$\pi_i^t(p_i^t, p_j^t) = \begin{cases} p_i^t a^t & \text{if } p_i^t < p_j^t \\ p_i^t a^t / 2 & \text{if } p_i^t = p_j^t \\ 0 & \text{if } p_i^t > p_j^t \end{cases}$$

As usual, a history at t is a sequence of the form $((a^1, p_1^1, p_2^1), \dots, (a^{t-1}, p_1^{t-1}, p_2^{t-1}), a^t)$, and a strategy s_i is a sequence (s_i^1, s_i^2, \dots) , where s_i^t prescribes a price after each

¹¹As we discuss below, the earlier literature sometimes used models of the conjectural variations style, but these models were somewhat unsatisfactory or even confusing.

possible history at t . A pair of strategies $s = (s_i, s_j)$ induces a probability distribution over infinite histories. Let E denote the expectation with respect to this distribution. Firm i 's payoff from the profile s is

$$\Pi_i(s) = E\left[\sum_{t=1}^{\infty} \delta^{t-1} \pi_i^t(p_i^t, p_j^t)\right]$$

where $\delta \in (0, 1)$ is the discount factor. The solution concept is SPE. There are of course many SPE's. This model focuses on a symmetric SPE that maximizes the firms' payoffs over the set of all SPE's.

Proposition 4.1: *There exists a symmetric SPE which maximizes the total payoff. Along its path, $p_i^t = p_j^t = p(a^t)$, where*

$$\begin{aligned} p(L) = p(H) = 1 & \quad \text{for } \delta > \frac{H}{(1+w)H+(1-w)L} \\ p(L) = 1, p(H) = \frac{\delta(1-w)L}{H[1-\delta(1+w)]} & \quad \text{for } \frac{H}{(1+w)H+(1-w)L} \geq \delta \geq \frac{1}{2} \\ p(H) = p(L) = 0 & \quad \text{for } \delta < 1/2 \end{aligned}$$

The proof is relegated to the appendix. The interesting part of this result obtains for the middle range of δ 's. Over this range, the equilibrium price during the high-demand state, $p(H)$, is lower than the monopoly price of 1. (For this range of δ , $p(H) = \delta(1-w)L/H[1-\delta(1+w)] < 1$). On the other hand, in the low-demand state, the equilibrium price achieves the monopoly price of 1. Rotemberg and Saloner refer to this result of lower prices during periods with higher demand as "price wars during booms." They argue that this is consistent with evidence on the behavior of oligopolistic industries over different phases of the business cycle. The intuition behind this result may be understood in the following general terms. A firm is willing to collude if the losses from future punishment outweigh the firm's immediate gain from deviation. In this model, owing to the independence of the shocks, the future losses are the same in "booms" and "busts." The immediate gains from defection at a given price, however, are obviously higher in booms. Therefore, when δ is not too large, it is impossible to sustain the monopoly price in a boom. To sustain collusion in a boom, it is necessary to reduce the temptation to deviate by colluding at a lower price.

DISCUSSION

Let us highlight three points arising from this analysis. The first point concerns the substance. Rotemberg and Saloner develop the general point that the pattern of collusion and the dynamics of demand are related in a predictable way.

Their analysis also uncovers the specific result that collusive prices are lower in high-demand states. This result derives to some extent from the assumption that demand shocks exhibit serial independence, and subsequent work has modified this assumption and reconsidered the relationship between demand levels and collusive prices. One finding is that collusion is easier to maintain (i.e., collusive prices tend to be higher) when future demand growth is expected to be large. The modified model can be also applied to markets with seasonal (as opposed to business-cycle) demand fluctuations, where the considerations of the oligopolists might be more transparent due to the greater predictability of the fluctuations. Indeed, recent empirical efforts offer evidence that is supportive of this hypothesis.¹² In any case, whether we consider the Rotemberg-Saloner model as is or one of its modified versions, the general approach suggested by this framework reveals considerations that plausibly influence the actual relationship between collusive prices and demand.

The second point concerns the essential role of the model. It should be noted that the repeated game model here is not incidental to the analysis. The main result of the analysis derives explicitly from the trade-off that the firm faces in a repeated game between the short-term benefit from undercutting the rival's price and the long-term cost of the consequent punishment. The result would seem somewhat counter-intuitive if one ignored the game theoretic reasoning and instead considered the situation using the standard price theoretic paradigms of monopoly and perfect competition. Of course, those who have the repeated game reasoning seated in the back of their minds can intuit through this argument easily, and may come to think that the formal model is superfluous. But then one has to have this reasoning already in the back of one's mind and to associate it with oligopolistic collusion.

The third point calls attention to one of the important strengths of this framework. The analysis illustrates clearly the flexibility with which the basic model can be adapted to incorporate alternative assumptions on the environment (in Section 6 this point will be illustrated further by another application that incorporates imperfect information into this framework). It would be difficult or

¹²Haltiwanger and Harrington (1991) hypothesize that demand rises and then falls as part of a deterministic cycle, and they find that collusive prices are higher when demand is rising. This model is well-suited for markets that are subject to seasonal demand movements, and Borenstein and Shephard (1996) report evidence of pricing in the retail gasoline market that supports the main finding. Bagwell and Staiger (1997) hypothesize that the demand growth rate follows a Markov process, so that demand movements are both stochastic and persistent, and they find that collusive prices are higher in fast-growth (i.e., boom) phases.

even impossible to meaningfully incorporate such features into the conjectural variations paradigm (see Section 7 below).

We chose to devote much of the above presentation to a rather specific model. It would therefore be useful in closing to take a broader view and call attention to two fundamental insights of the repeated games literature that contain important lessons for oligopoly theory. The first insight is that a collusive outcome, which serves the firms better than the one-shot equilibrium, can be sustained in the interaction of fully rational competitors (in the sense formalized by the notion of SPE). The second is that repetition of the one-shot equilibrium is a robust outcome of such interaction: it is always a SPE and, in many interesting scenarios (e.g., finite horizon, high degree of impatience, short memory), it might even emerge as the unique one. The first insight provides a clear theory of collusion which can identify some key factors that facilitate collusion. The second shows the relevance of the static oligopoly models and the insights they generate: their equilibria continue to have robust and sometimes unique existence in a much richer dynamic environment. Of course, the consideration of the collusive and non-collusive outcomes predated the more recent analyses of the repeated game model. The important contribution of the repeated game framework is in establishing the validity of these as outcomes of rational and far sighted competition that takes place over time.

5. Sales: an application of mixed-strategies equilibrium

In many retail markets prices fluctuate constantly and substantially. At any given point in time, some firms may offer a “regular” price, while other firms temporarily cut prices and offer “sales” or price promotions. The frequency and the significance of these price movements make it hard to believe that they mirror changes in the underlying demand and cost conditions.

The earlier literature had largely ignored these phenomena. Sales did not fit well into the existing price theoretic paradigm, and as a result this practice may have been viewed as reflecting irrational behavior that was better suited for psychological study than for economic analysis.

The game theoretic notion of a mixed-strategy equilibrium presents an alternative view whereby the ubiquitous phenomenon of sales can be interpreted as a stable outcome of interaction among rational players. We develop this argument

with Varian’s (1980) model of retail pricing.¹³ This model highlights a tension that firms face between the desire to price high and profit on consumers that are poorly informed as to the available prices and the desire to price low and compete for consumers that are well informed of prices. This tension is resolved in equilibrium when firms’ prices are determined by a mixed strategy. Varian’s theory thus predicts that firms will offer sales on a random basis, where the event of a sale is associated with the realization of a low price from the equilibrium mixed strategy. A dynamic interpretation of this theory further implies that different firms will offer sales at different points in time.

A notable feature of this theory is that it takes the random behavior of the mixed-strategy equilibrium quite seriously and uses it to directly explain the randomness in prices observed in real markets. At the same time, the mixed-strategy approach has a well-known potential drawback, associated with the literal interpretation of the assumption that each firm selects its price in a random manner. This difficulty is often addressed with reference to Harsanyi’s (1973) idea that a featured mixed-strategy equilibrium for a given game can be re-interpreted as a pure-strategy Nash equilibrium for a “nearby” game of incomplete information. However, there is of course the question of the plausibility of the nearby game for the application of interest. With these concerns in mind, we develop as well an explicit and plausible “purification” of the mixed-strategy equilibrium for Varian’s pricing game. This analysis suggests that the random pattern of sales also can be understood as reflecting small private cost shocks that vary across firms and time.

AN EQUILIBRIUM THEORY OF SALES

We begin with the basic assumptions of the model. A set of $N \geq 2$ symmetric firms supplies a homogeneous good at unit cost c to a consumer population of unit mass. Each consumer demands one unit of the good, and the good provides a gross utility of v , where $v > c \geq 0$. There are two kinds of consumers. A fraction $I \in (0, 1)$ of consumers are informed about prices and hence purchase from the firm with the lowest price; if more than one low-priced firm exists, these consumers divide their purchases evenly between the low-priced firms. The complementary fraction $U = 1 - I$ of consumers are uninformed about prices and hence pick firms from which to purchase at random. Given these assumptions, we define a simultaneous-move game played by N firms as follows. A pure strategy for any firm i is a price $p_i \in [c, v]$. Letting p_{-i} denote the $(N - 1)$ -tuple of prices selected

¹³Related models were also explored by Shilony (1977) and Rosenthal (1980).

by firms other than firm i , the profit to firm i is defined as:

$$\Pi_i(p_i, p_{-i}) = \begin{cases} [p_i - c]U/N & \text{if } p_i > \min_{j \neq i} p_j \\ [p_i - c](U/N + I/k) & \text{if } p_i \leq \min_{j \neq i} p_j \text{ \& \; } \|\{j : p_j = p_i\}\| = k - 1 \end{cases} \quad (5.1)$$

A mixed strategy for a firm is a distribution function defined over $[c, v]$. If firm i employs the mixed strategy F_i and the strategies of its rivals are represented with the vector F_{-i} , then the expected profit to firm i is:

$$E_i(F_i, F_{-i}) = \int_c^v \cdots \int_c^v \Pi_i(p_i, p_{-i}) dF_1 \dots dF_N$$

For this game, a price vector $\{p_1, \dots, p_N\}$ forms a Nash equilibrium in pure strategies if, for every firm i and every $p'_i \in [c, v]$, we have $\Pi_i(p_i, p_{-i}) \geq \Pi_i(p'_i, p_{-i})$. Allowing also for mixed strategies, the distributions (F_1, \dots, F_N) form a Nash equilibrium if, for every firm i and every distribution function F'_i , we have that $E_i(F_i, F_{-i}) \geq E_i(F'_i, F_{-i})$. A symmetric Nash equilibrium is then a Nash equilibrium (F_1, \dots, F_N) satisfying $F_i = F$, for all $i = 1, \dots, N$. Let $\bar{p}(F)$ and $\underline{p}(F)$ denote the two endpoints of the support of F ; i.e., $\bar{p}(F) = \inf\{p : F(p) = 1\}$ and $\underline{p}(F) = \sup\{p : F(p) = 0\}$. We may now present the main finding:

Proposition 5.1: (A) *There does not exist a pure-strategy Nash equilibrium.*

(B) *There exists a unique symmetric Nash equilibrium F . It satisfies:*

- (i) $\bar{p}(F) = v$;
- (ii) $[\underline{p}(F) - c](U/N + I) = [v - c](U/N)$
- (iii) $[p - c](U/N + (1 - F(p))^{N-1}I) = [v - c](U/N)$ for every $p \in [\underline{p}(F), \bar{p}(F)]$.

The proof is relegated to the appendix. The proposition reflects the natural economic tension that each firm faces between the incentive to cut price—offer a “sale”—in order to increase its chances of winning the informed consumers and the incentive to raise its price in order to better profit on its captive stock of uninformed consumers. This tension precludes the existence of a pure-strategy equilibrium, since the presence of informed consumers induces a firm to undercut its rivals when price exceeds marginal cost, while the presence of uninformed consumers induces a firm to raise its price when price equals marginal cost. In the mixed-strategy equilibrium, these competing incentives are resolved when firms select prices in a random manner, with some firms offering sales and other firms electing to post higher prices.

Now, while the specific predictions of the model seem to accord with casual observations and also with formal empirical studies,¹⁴ the very literal and direct use of mixed strategies to explain price fluctuations raises some questions of interpretation. Do firms really select prices in a random manner? Correspondingly, are firms really indifferent over all prices in a certain range? And, if so, what compels a firm to draw its price from the specific equilibrium distribution? To address these questions, we develop next a purification argument that applies to this game.

PURIFICATION

We apply here Harsanyi's (1973) idea to interpret the mixed-strategy equilibrium of this game as an approximation to a pure-strategy equilibrium of a nearby game with some uncertainty over rival's costs. The uncertainty ensures that a firm is never quite sure as to the actual prices that rivals will select, and so incomplete information plays a role analogous to randomization in the mixed-strategy equilibrium presented above.

Consider now an incomplete-information version of the above game in which the firms' cost functions are private information. Firm i is of type $t_i \in [0, 1]$. A firm knows its own type but it does not know the types of the other firms. It believes that the types of the others are realizations of i.i.d random variables with uniform distribution over $[0, 1]$. The firm's type determines its cost function: firm i of type t_i has cost $c(t_i)$, where the function c is differentiable and strictly increasing and $0 < c(0) < c(1) < v$. As before, the firms simultaneously choose prices and receive the corresponding market shares and profits. Thus, this model is a standard Bayesian game with type spaces $[0, 1]$ and uniformly distributed beliefs. Notice that the uniform distribution of the beliefs is without loss of generality, since any differentiable distribution of costs can still be obtained by the appropriate choice of the cost function c . In this game, a pure strategy for any firm i is a function, $P_i(t_i)$, that maps $[0, 1]$ into $[c(0), v]$. Given a strategy profile $[P_1, \dots, P_N]$, let P_{-i} denote the strategies of the firms other than i and let $P_{-i}(t_{-i})$ denote the vector of prices prescribed by these strategies when these firms' types are given by the $(N - 1)$ -tuple t_{-i} . The profit of firm i of type t_i that charges p_i when its rivals are of types t_{-i} is $\Pi_i(p_i, P_{-i}(t_{-i}), t_i)$ where Π_i is given by (5.1) in which c is replaced with $c(t_i)$. The profile $[P_1, \dots, P_N]$ is a Nash equilibrium if, for

¹⁴For example, Villas-Boas (1995) uses price data on coffee and saltine cracker products, and argues that the pricing patterns observed for these products are consistent with the Varian model.

all i and all t ,

$$P_i(t_i) \in \arg \max_{p_i} E_{t_{-i}}[\Pi_i(p_i, P_{-i}(t_{-i}), t_i)]$$

A symmetric Nash equilibrium is such that $P_i(t_i) = P(t_i)$ for all i and t_i .

The distribution of prices induced by a strictly increasing strategy P is given by $G(x) = \text{Prob}\{t \mid P(t) \leq x\} = P^{-1}(x)$.

Proposition 5.2: (i) *In the incomplete-information game, there exists a pure-strategy and strict Nash equilibrium, P .*

(ii) *Given a constant $c \in (0, v)$, for any $\varepsilon > 0$, there exists $\delta > 0$ such that, if $|c(t) - c| < \delta$ for all t , then the distribution of prices induced by P , $P^{-1}(x)$, is ε close to F_c , the mixed-strategy equilibrium of the complete-information game with common per-unit cost c .*

The proof is relegated to the appendix. In other words, the pure-strategy Nash equilibrium that arises in the incomplete-information game when the costs are near c for all t generates approximately the same distribution over prices as occurs in the mixed-strategy equilibrium of the complete-information game with common cost c . The mixed-strategy equilibrium for Varian's game can thus be interpreted as describing a pure-strategy equilibrium for a market environment in which firms acquire cost characteristics that may differ slightly and are privately known.

DISCUSSION

The important substantive message of this theory is that price fluctuations which take the form of sales and promotions are largely explained as a consequence of straightforward price competition in the presence of buyers with varying degrees of information, rather than by some significant exogenous randomness in the basic data. When we adopt the interpretation of the purified version, the intuition can be described as follows. When there is a mix of better informed and less informed consumers, there are conflicting incentives to price high and low as explained above. Price competition arbitrates the low and high prices to the point where they are nearly equally profitable. In such a situation, relatively small differences in the firms' profit functions, such as those caused by small cost shocks, can yield large price movements. What this insight means for the empiricist is that the relevant data for understanding such markets concerns perhaps the nature of consumer information to a larger extent than it concerns technological or taste factors.

The predictions of this model appear even more compelling, when one looks at the immediate dynamic extension. Like in any other static oligopoly game, the featured one-shot equilibrium corresponds to the non-collusive SPE of the repeated version of this game. We mention this obvious point specifically, since when the mixed strategy is played repeatedly, different firms can be expected to offer sales in different periods, which is a prediction that seems consistent with casual observation. In the dynamic purified version, each firm is privately informed of its current cost at the start of each period, and the cost shocks are assumed independent across time. The periodic cost shock might reflect, for example, firm-specific data like the level of the firm's inventories and the extent to which it is pressed for storage space.

Notice that the use of the game theoretic model here goes beyond a formal exposition of some natural intuition. In the absence of systematic equilibrium analysis and the concept of equilibrium in mixed strategies, it would be difficult or even impossible to come up with this explanation. Only once the result is obtained, it becomes possible to understand it intuitively.

6. On the contribution of industrial organization theory to game theory

Although industrial organization theory has been mainly a user of concepts and ideas which had been generated by game theory without an explicit industrial organization motivation, the relationship has not been totally one-sided. There are specific ideas that grew out of problems in industrial organization that gained independent importance as game theoretic topics in their own right. In what follows, we describe in detail one idea of this nature.

REPEATED GAMES WITH IMPERFECT MONITORING – “PRICE WARS”

The development of this model was motivated by the observation that some oligopolistic industries experience spells of relatively high prices, which seem to result from implicit collusion, interrupted by spells of more aggressive price competition, referred to as “price wars.”¹⁵ A somewhat trivial theory could point out that this is consistent with the paths of certain SPE in a repeated oligopoly game. In such an equilibrium the firms coordinate for a while on collusive behavior (high prices or small quantities), then switch for a while to the one-shot equilibrium,

¹⁵Porter (1983) studies this pattern of behavior among firms in the U.S. railroad industry.

and so on. What makes this theory rather unconvincing is that the alternation between collusion and price warfare is an artificial construct which does not reflect some more intuitive considerations. Moreover, there are simpler equilibria which Pareto dominate an equilibrium of this form. A more interesting theory for the instability of oligopolistic collusion was suggested by Rotemberg and Saloner (1986), as reviewed above, but this explanation is confined to price wars triggered by foreseen variations in demand.

The theory suggested by Green and Porter (1984) views the oligopolistic interaction as a repeated game with imperfect public information. In the repeated Cournot version of this approach (which is the version analyzed by Green and Porter), the firms simultaneously choose outputs in each period, and the price is a function of the aggregate output and some random demand shock which is unobservable to the firms. The firms observe the price but cannot observe their rivals' outputs; consequently, a firm cannot tell whether a low price is the result of a bad demand shock or a high output by some rival. While the firms would like to collude on producing smaller outputs than those entailed by the static Cournot equilibrium, in this imperfect-monitoring environment it is impossible to sustain uninterrupted collusion. Intuitively, low prices cannot always go unpunished, since then firms would be induced to deviate from the collusive behavior. But this implies that collusion must sometimes break down into a "price war" along the equilibrium path. Reasoning in this way, Green and Porter constructed equilibria which exhibit on their paths spells of collusive behavior interrupted by blocks of time (following bad demand shocks) during which the firms revert to playing the Cournot equilibrium of the one-shot game.

It is somewhat easier to illustrate this point using a simple version of a repeated Bertrand duopoly game.¹⁶ Two firms produce a homogenous product at zero cost. The demand depends on the state of nature: with probability α there is no demand, and with probability $(1 - \alpha)$ the demand is a simple step function

$$Q(p) = \begin{cases} 2 & \text{if } p \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

The firms simultaneously choose prices $p_i \in [0, 1]$. If the prices are equal, the firms share the demand equally; otherwise, the low-price firm gets the entire

¹⁶The discussion here is influenced by Tirole's (1988) presentation.

demand. The payoffs of the firms in the high-demand state are

$$\pi_i(p_i, p_j) = \begin{cases} 2p_i & \text{if } p_i < p_j \\ p_i & \text{if } p_i = p_j \\ 0 & \text{if } p_i > p_j \end{cases}$$

and in the zero-demand case the payoffs to firms are zero. The firms do not observe the realization of the demand directly, but only their own shares. So, if both charged price p and the demand was high, these facts are public information. Otherwise, the only public information is that no such event occurred.

In the repeated game version, this interaction is repeated in each period $t = 1, 2, \dots$. The firms' payoffs are the discounted sums of their profits with a common discount factor δ . In addition, assume that at the end of each period t the firms commonly observe a realization of a random variable, x^t , distributed uniformly over $[0, 1]$ and independently across periods. This variable is a mere “sunspot” which does not affect the demand or any other of the “real” magnitudes, but the possibility to condition on it enriches the set of strategies in a way that simplifies the analysis. A public history of the game at t is a sequence $h^t = (a^1, \dots, a^{t-1})$, where either (i) $a^r = (p, x)$ which means that, in period r , demand was high, both prices were equal to $p > 0$ and the realization x was observed, or (ii) $a^r = (\sim, x)$ which means that either both prices were 0 or at least one of the firms sold nothing at r . A strategy for firm i prescribes a price choice for each period t after any possible public history. A sequential equilibrium (SE) is a pair of strategies (which depend on public histories) such that i 's strategy is best response to j 's strategy, $i \neq j = 1, 2$, after any public history.

In the one-shot game, the only equilibrium is the Bertrand Equilibrium: $p_i = 0$, $i = 1, 2$. Indefinite repetition of this equilibrium is of course a SE in the repeated game. If the demand state became observable at the end of each period, then provided that δ is sufficiently large, the repeated game would have a perfectly collusive SPE in which $p_1 = p_2 = 1$ in perpetuity. It is immediate to see that, in the present model, there is no perfectly collusive SE. If there were such a SE, in which the firms always choose $p_i = 1$ along the path, it would have to be that firm i continues to choose $p_i = 1$ after periods in which it does not get any demand. But, then it would be profitable for firm j to undercut i 's price.

The interesting observation from the viewpoint of oligopoly theory is that there are equilibria which exhibit some degree of collusion and that, due to the impossibility of perfect collusion, such equilibria must involve some sort of “price warfare” on their path. Green and Porter identified a class of such equilibria that

alternate along their path between a collusive phase and a punishment phase. In the present version of the model, these equilibria are described in the following manner. In the collusive phase the firms charge $p_i = 1$, and in the punishment phase they charge $p_i = 0$. The transition between the phases is then characterized by a nonnegative integer T and a number $\beta \in [0, 1]$. The punishment phase is triggered at some period t by a “bad” public observation of the form $a^{t-1} = (\sim, x)$, where $x < \beta$; the collusive phase is restarted after T periods of punishment.

To construct such a SE, let T and β be as above. Define the set of (“good”) histories $G(T)$ to consist of: (i) the empty history; (ii) histories that end with $(1, x)$ for any x ; (iii) histories such that since the beginning or since the last observation of the form $(1, x)$ there has been a block of exactly $k(T + 1)$ observations of the form (\sim, x) , where k is a natural number. Define the strategy $f_{T,\beta}$ as follows

$$f_{T,\beta}(h) = \begin{cases} 1 & \text{if } h \in G(T) \\ 1 & \text{if } h = (h', (\sim, x)) \text{ where } h' \in G(T) \text{ \& } x \geq \beta \\ 0 & \text{otherwise} \end{cases}$$

Thus, the occurrence of the bad demand state does not always trigger the punishment, but only when $x < \beta$. Suppose that both firms play this strategy and let $V_{T,\beta}$ denote the expected discounted payoff for a firm calculated at the beginning of a period, after a history that belongs to $G(T)$:

$$V_{T,\beta} = 1 - \alpha + (1 - \alpha\beta)\delta V_{T,\beta} + \alpha\beta\delta^{T+1}V_{T,\beta} \quad (6.1)$$

The RHS captures the fact that, when both follow this strategy, with probability $\alpha\beta$, there is no demand and $x < \beta$, so the interaction will switch to the punishment phase for T periods; and with probability $1 - \alpha\beta = \alpha(1 - \beta) + (1 - \alpha)$, there is either no demand and $x \geq \beta$ or there is high demand, in which cases firms will continue colluding in the following period. Each firm gets a unit profit in the current period with probability $1 - \alpha$. Rearrangement of the above gives

$$V_{T,\beta} = (1 - \alpha)/[1 - \delta + \alpha\beta(\delta - \delta^{T+1})] \quad (6.2)$$

To verify that these strategies constitute an equilibrium, it is enough to check that there is no profitable single-period deviation after histories such that $f_{T,\beta}(h) = 1$ (since $p_1 = p_2 = 0$ is a Nash Equilibrium of the one-shot game there is clearly no incentive to deviate after other histories). Thus, the equilibrium condition is

$$2(1 - \alpha) + (1 - \beta)\delta V_{T,\beta} + \beta\delta^{T+1}V_{T,\beta} \leq V_{T,\beta} \quad (6.3)$$

The LHS captures the value of a single-period deviation. The payoff 2 is the supremum over the immediate payoffs that a firm can get by undercutting its rival's price. Since the deviation yields public information of the form (\sim, x) , the continuation will be determined by the size of x : with probability $(1 - \beta)$, $x \geq \beta$ so the collusion will continue in the next period yielding the value $\delta V_{T,\beta}$; and with probability β , $x < \beta$ so the T -periods punishment phase begins yielding the value $\delta^{T+1} V_{T,\beta}$, associated with the renewed collusion after T periods. Rearrange (6.3) to get

$$V_{T,\beta} \geq 2(1 - \alpha)/[1 - \delta + \beta(\delta - \delta^{T+1})] \quad (6.4)$$

Proposition 6.1: (i) *There exists an equilibrium of this form (with possibly infinite T), iff*

$$\alpha \leq 1 - \frac{1}{2\delta} \quad (6.5)$$

(ii) *For any α and δ satisfying (6.5), let $T(\alpha, \delta) = \min\{T \mid (1 - \delta)/(1 - 2\alpha)\delta(1 - \delta^T) \leq 1\}$.*

$$\arg \max_{T,\beta} [V_{T,\beta} \text{ s.t. (6.3)}] = \{(T, \beta) \mid T \geq T(\alpha, \delta) \text{ and } \beta = \frac{1 - \delta}{(1 - 2\alpha)\delta(1 - \delta^T)}\}$$

Proof. (i) Substitute from (6.2) to the LHS of (6.4) to get that a (T, β) equilibrium exists iff

$$(1 - \alpha)/[1 - \delta + \alpha\beta(\delta - \delta^{T+1})] \geq 2(1 - \alpha)/[1 - \delta + \beta(\delta - \delta^{T+1})] \quad (6.6)$$

Rearrangement yields

$$\alpha \leq \frac{1}{2} - \frac{1 - \delta}{2\beta(\delta - \delta^{T+1})} \quad (6.7)$$

Since the RHS increases with T and β , if this inequality holds for some T and β , it must hold for $T = \infty$ and $\beta = 1$. Thus, there are some T and β for which (6.7) holds iff (6.5) holds.

(ii) Let (T, β) be an equilibrium configuration that maximizes $V_{T,\beta}$. For $T' \geq T$, define $\beta' = \beta(1 - \delta^T)/(1 - \delta^{T'})$ and observe that (T', β') is also an equilibrium configuration that maximizes $V_{T,\beta}$. To see this, note first that $\beta' \leq \beta \leq 1$. Second, since (6.6) holds for (T, β) , it holds for (T', β') as the denominator remains the same and this implies that it is an equilibrium. Third, $V_{T',\beta'} = V_{T,\beta}$, since (6.2) gets the same value with (T, β) and (T', β') . Thus, in particular, an equilibrium with $T = \infty$ is always among the maximizers of $V_{T,\beta}$.

Next, observe from (6.2) that $V_{\infty,\beta}$ is decreasing in β , so that $V_{\infty,\beta}$ is maximized at the minimal β that satisfies (6.7) with $T = \infty$, which is $\beta_{\infty} = (1 - \delta)/(1 - 2\alpha)\delta$. Now, inspection of (6.2) implies that, for an equilibrium with $T < \infty$ to be also a maximizer of $V_{T,\beta}$, it has to be that the β of this equilibrium satisfies $\beta = \beta_{\infty}/(1 - \delta^T)$. This is possible only for T 's such that $\beta_{\infty}/(1 - \delta^T) \equiv (1 - \delta)/(1 - 2\alpha)\delta(1 - \delta^T) \leq 1$, i.e., only for $T \geq T(\alpha, \delta)$. ■

This proposition shows that, for some range of the parameters, there are equilibria which exhibit the sought after form of behavior: spells of collusion interrupted by price wars. Part (ii) shows that there are such equilibria among the optimal ones in the T, β -class. Moreover, a result due to Abreu, Pearce and Stacchetti (1986) implies that, in this model, the optimal equilibria in the T, β -class are also optimal among all symmetric equilibria (not just optimal in the T, β -class). So the sought after alternation between collusion and price wars on the equilibrium path emerges, even when we insist on optimal symmetric equilibria.

This observation is somewhat qualified by the fact that the $T = \infty$ equilibrium, which does not alternate between the two regimes, is always among the optimal equilibria as well. The present model however is rather special. In particular, it has the property that the worst punishments that the parties can inflict on one another coincide with the Nash equilibrium of the one-shot game. Abreu, Pearce and Stacchetti (1986) show in a more general model that there is a joint profit maximizing symmetric equilibrium that starts with playing the most collusive outcome and then switches into the worst sequential equilibrium. However, in models such as the repeated Cournot game, in which the one-shot equilibrium does not coincide with the worst punishment, the worst sequential equilibrium itself would involve alternation between the collusive outcome and another “punishing” outcome. So the behavior along the path of the optimal symmetric equilibrium would still have the appearance of collusion interrupted by price wars. The narrative that accompanies this equilibrium is less direct: the spells of collusion are in some sense rewards for sticking to the punishment and they are hence triggered by sufficiently bad public signals (low prices in the Cournot version) which confirm the firms’ adherence to the punishment. In this sense, the behavior captured by these equilibria is not as “natural” or “straightforward” as the behavior captured by the original Green-Porter equilibrium.

The relevance of imperfect monitoring for cartel instability gets an additional twist once asymmetric equilibria are considered. Fudenberg, Levine and Maskin (1993) show that, when the public signal satisfies a certain full dimensionality property and δ is sufficiently close to 1, there are (asymmetric) equilibrium out-

comes which are arbitrarily close to the Pareto efficient outcome. Thus, under such conditions, the extent of “price warfare” is insignificant along the path of the optimal equilibrium.

The imperfect-monitoring model of collusion contributes importantly to the understanding of cartel instability. In addition, while it obviously started from a clear industrial organization motivation, this model has also generated a research line in the theory of repeated games with unobservable moves that has assumed a life of its own and that continues to grow in directions that are now largely removed from the original motivation.¹⁷

7. An overview and assessment

Non-cooperative game theory has become the standard language and the main methodological framework of industrial organization. Before the onset of game theory, industrial organization was not without analytical methodology — the highly developed methodology of price theory served industrial organization as well. But traditional price theory addresses effectively only situations of perfect competition or pure monopoly, while industrial organization theory emphasizes the spectrum that lies between these two extremes: the study of issues like collusion and predation simply requires an oligopoly model. This gap was filled by verbal theorizing and an array of semi-formal and formal models. The formal models included game models like those of Cournot, Bertrand and Stackelberg, as well as non-game models with strategic flavor such as the conjectural variations and the contestable market models. Before proceeding with the discussion, we pause here to describe the conjectural variations model,¹⁸ which is an important representative of the formal pre-game theoretic framework.

A pre-game theoretic model: the conjectural variations model

This model attempts to capture within an atemporal framework both the actual actions of the oligopolists and their responses to each other’s choices. In

¹⁷Repeated game models with imperfect monitoring had been considered somewhat earlier by Rubinstein (1979) and Radner (1981), who analyzed repeated Principal-Agent relationships. Besides the different basic game, these contributions also differ in their solution concepts (Stackelberg and Epsilon-Nash respectively) and their method of evaluating payoff streams (limit-of-the-means criterion). It seems however that, due to these differences or other reasons, the Green-Porter article has been more influential in terms of stimulating the literature.

¹⁸For a traditional description of this model, see Fellner (1949); for a modern view from the game theoretic perspective, see Friedman (1990).

the quantity-competition duopoly version, two firms, 1 and 2, produce outputs, q_1 and q_2 , which determine the price, $P(q_1 + q_2)$, and hence the profits, $\pi_i(q_i, q_j) = q_i P(q_1 + q_2) - c_i(q_i)$. Firm i holds a conjecture, q_i^C , regarding j 's output. In the conjectural variations framework, this conjecture may in fact depend on firm i 's own choice, q_i ; i.e., $q_i^C = v_i(q_i)$. An equilibrium is then a pair of outputs, q_i^* , $i = 1, 2$, such that

$$q_i^* = \arg \max_{q_i} \pi_i[q_i, v_i(q_i)] \quad \text{and} \quad v_i(q_i^*) = q_j^* \quad i = 1, 2.$$

Thus, each firm maximizes its profit under the conjecture that its rival's output will vary with its own choice, and the conjecture is not contradicted at the equilibrium.

In many applications, the v_i 's were assumed linear in the q_i 's (at least in the neighborhood of the solution) with slope v . Under this assumption, the parameter v indexes the equilibrium level of collusion: with $v = 1, 0$ and -1 , the equilibrium outcome coincides with the joint monopoly outcome, the Cournot equilibrium outcome and the perfectly competitive outcome, respectively, and with other v 's in this range the outcome falls between those mentioned.

Notice that, unlike non-cooperative game theoretic models, this model remains vague about the order of moves. In fact, if we tried to fit this model with an extensive form, it would have to be such that each firm believes that it is moving ahead of the other firm. This is possible only if the firms hold inconsistent beliefs (say, different prior beliefs over the moves of nature who chooses the actual sequence of moves). But obviously this model was not meant to capture behavior under a special form of inconsistent beliefs. It is probably more appropriate to think of it as a reduced form of an underlying dynamic interaction that is left unmodelled.

It is important to note that, in the pre-game theoretic literature, the conjectural variations model was not viewed as different in principle from the game theoretic models. This is because the game models were viewed then somewhat differently than they are viewed now. They were not seen as specific applications of a very deep and encompassing theory, the Nash equilibrium theory, but rather as isolated specific models using a somewhat ad-hoc solution concept. In fact, the Nash equilibria of these models were often viewed as a special case of the conjectural variations model and were often referred to as the "zero-conjectural-variations" case.

Having mentioned the theoretical background against which the game theoretic models were introduced, let us try to assess briefly the contribution of this

change in the theoretical framework of industrial organization. The following points discuss some aspects of this contribution both to the expositional role of the theory and to its substance.

Game theory as a language. The first contribution of game theory to industrial organization is the introduction of a language: models are described in an accurate and economical way using standard familiar formats, and basic non-cooperative solution concepts are commonly employed. One clear benefit of this standardization is improved accessibility: a formal political theorist or a mathematician can access relatively easily the writings in modern industrial organization theory, without a long introduction to the specific culture of the field. This requires, of course, some basic familiarity with the language of game theory. But the pre-game theoretic literature also had by and large a language of its own – it was just not as universal and as efficient as game theory.

To appreciate the contribution of game theory simply as a language, one merely has to read through some of the presentations of formal models in the pre-game theoretic literature. The above description of the conjectural variations model already benefited from the game theoretic language and perhaps does not reflect appropriately the ambiguity that often surrounded the presentation of this model. The ambiguity that naturally results from the timeless unmodelled dynamics was further exacerbated in many cases by presentations that described the central assumptions of the model only in terms of derivatives (perhaps to avoid the embarrassment of describing the inconsistent beliefs explicitly).

Game theory has the additional virtue of being a more flexible language, in the sense that it allows consideration of situations involving dynamic interaction and imperfect information regarding behavior and the environment. The flexibility of the game theoretic framework in these respects derives perhaps from the rather primitive structure of non-cooperative game models, which set forth all actions and their timing. To be sure, the consideration of dynamic interaction and imperfect information complicates the models and raises additional conceptual problems (say, about the appropriate solution concept). Nevertheless, game theory does provide a systematic and formal language with which to explore these considerations - something that was not available in the pre-game theoretic literature. For example, in the conjectural variations model of collusion, it is not even clear how to begin thinking about the consequences of secret price cutting for collusive conduct. This model is not suitable for such an analysis, since it fails to describe the information that firms possess and the timing of their actions. Furthermore, there is no obvious way to add these dimensions. By contrast, such

an analysis is natural in the context of the repeated game model of collusion, as discussed in the Green-Porter model reviewed in Section 6. This flexibility is an important asset of the game theoretic framework.

Game theory as a discipline. Related to its role as a language, game theory also imposes a discipline on modeling. First, a non-cooperative game model requires the analyst to specify precisely the actions available to players, the timing of the actions and the information held by players. This forces modelers to face their assumptions and hence to question them. For example, the repeated game collusion models of Sections 4 and 6 are very specific about the behavior and the information of firms. By contrast, in the conjectural variations model, no explicit assumptions on behavior are presented, so that one can judge the model only by the postulated outcome.

Second, with the game theoretic framework, results have to satisfy the requirements of known solution concepts, usually Nash equilibrium and its refinements. This forces arguments to be complete in the way dictated by these solutions. The brief review of the development of the literature on entry deterrence at the end of Section 2 illustrates this point. The argument on the role of investment and price in deterrence became complete, only after the situation was described as a simple two-stage game and analyzed with the appropriate solution concept of SPE.

The imposition of the game theoretic discipline has some drawbacks as well. First, it naturally constrains the range of ideas that can be expressed and inhibits researchers from venturing beyond its boundaries. Second, careless use of game theory may lead to substantial misconceptions. The Nash equilibrium concept is not always compelling. Standard game theoretic models often presume high degrees of rationality and knowledge on the part of the players, and the full force of such assumptions is often not acknowledged in applications. Third, there is a sense in which non-cooperative game models require the modeler to specify too much. The game theoretic collusion models specify whether the firms move simultaneously or alternately, what exactly they observe, and so on. These features are normally not known to observers and natural intuition suggests that they should not be too relevant. However, they have to be specified and even worse they often matter a great deal for the analysis. If one were very confident about the accuracy of the description and the validity of the solution concept, the important role of the fine details of the model might provide important insights. But since the models are often viewed as a rather rough sketch of what goes on, the sensitivity of predicted outcomes to modeling details is bothersome. In contrast, the conjectural variations model summarizes complicated interaction simply, without

spending much effort on the specification of artificial features, so that the fine details do not interfere with the overall picture.

The substantive impact of game theory. So industrial organization has a new language/discipline and perhaps a superior one to what it had before, but has this language generated new insights to the substance of industrial organization? By taking a very broad view, one might argue that it has not. Take oligopoly theory, for example. In the pre-game theoretic era, economists clearly recognized the potential inefficiency of oligopolistic competition, the forces that work to induce and diffuse collusion and the possibility that different degrees of collusion could be sustained by threats of retaliation. In some sense, this is what is known now, too.

But a closer look reveals quite a few new specific insights. In fact, each of the previous sections described what we believe to be a new insight, and we attempted to identify the crucial role of the game theoretic framework in reaching these insights. For example, the idea that export subsidies can play a strategic role that might rationalize their use would be difficult to conceive without the game theoretic framework. In fact, it runs contrary to intuition based on standard price theory. Similarly, in the absence of this framework, it would be hard to conceive of the idea that random pricing in the form of sales is a robust phenomenon which derives from the heterogeneous price information (or search costs) that different segments of the consumer population enjoy.

But the fact that a certain relationship that exists within the model can be interpreted in terms of the underlying context, and is thus regarded as an insight, does not necessarily mean that it truly offers a better qualitative understanding of important aspects of actual market behavior. There remains the question of whether or not such an insight is more than a mere artifact of the model. For example, as discussed in section 2, the strategic explanation of export subsidization (taxation) might be questioned in light of the sensitivity of this explanation to modeling decisions (Cournot vs. Bertrand). All things considered, however, we believe that the insights described above identify qualitative forces that plausibly play important roles in actual markets.

At the same time, we also stress that these insights should not be taken too literally. For example, the model of Section 5 tells us that sales can be a stable phenomenon through which price competition manifests itself. The insight is that this phenomenon need not reflect some important instability in the technology or pattern of demand; rather, sales emerge naturally from price competition when consumers are heterogeneously informed. Of course, this is not to say that firms

know exactly some distribution and go through the precise equilibrium reasoning. The point is only that they have some rough idea that others are also pursuing these sales policies, and given this they have no clearly superior alternative than to also have sales in response to small private signals.

Let us accept then that many insights derived from the game theoretic approach offer a better qualitative understanding of important aspects of actual market behavior. We may still question the deeper significance of these insights. In particular, has the game theoretic framework delivered a new class of models that consistently facilitates better quantitative predictions than what would have been available in its absence? A serious attempt to discuss this question would take us well beyond the scope of this paper. Here, we note that important new empirical work in industrial organization makes extensive use of game theoretic models, but we also caution that there is as yet no simple basis from which to conclude that the game theoretic approach provides consistently superior quantitative predictions.

This inconclusive answer regarding the quantitative contribution of game theory does not imply that the usefulness of this framework for policy decisions is doubtful. Even if game theory has not produced a magic formula that would enable a regulator to make a definitive quantitative assessment as to the consequences of a proposed merger, this framework has enabled the regulators to think more thoroughly about the possible consequences of the merger. Likewise, it offers the regulator a deeper perspective on the issue of predation. To be sure, it does not offer a magic formula here either. But it makes it possible to have a more complete list of scenarios in which predation might be practiced and to use such arguments to justify intervention in situations that would not warrant intervention on the basis of simple price theoretic arguments.

8. Appendix

Proposition 3.1: (i) *There exists a separating intuitive equilibrium.*

(ii) *If $p_L^m \geq \underline{p}$, then any separating intuitive equilibrium, $\{P, E, b\}$, satisfies $\underline{p} = P(L) < P(H) = p_H^m$ and $E(P(L)) = 0 < 1 = E(P(H))$.*

If $p_L^m < \underline{p}$, then any separating intuitive equilibrium, $\{P, E, b\}$, satisfies $p_L^m = P(L) < P(H) = p_H^m$ and $E(P(L)) = 0 < 1 = E(P(H))$.

(iii) *If $p_L^m \geq \underline{p}$, and $b_L^0 \geq \tilde{b}_L$, then for every $p \in [\underline{p}, p_L^m]$, there exists an intuitive pooling equilibrium in which $P(L) = P(H) = p$.*

(iv) *In any intuitive pooling equilibrium, $P(L) = P(H) \in [\underline{p}, p_L^m]$ and $E(P(L)) =$*

0.

Proof. (i) For the case $p_L^m < \underline{p}$, define the triplet $\{P, E, b\}$ as follows: P is as in (ii) above, $E(p) = 1$ iff $p \neq p_L^m$, $b_L(p) = 0$ if $p \neq p_L^m$ and $b_L(p_L^m) = 1$. It is direct to verify that this triplet satisfies (E1)-E(4).

For the case $p_L^m \geq \underline{p}$, define the triplet $\{P, E, b\}$ as follows: P is as in (ii) above, $E(p) = 1$ iff $p \in (\underline{p}, \bar{p}]$, $b_L(p) = 0$ if $p \in (\underline{p}, \bar{p}]$ and $b_L(p) = 1$ otherwise. This triplet clearly satisfies (E1) when $t = H$ and when $t = L$ and $p_L^m = \underline{p}$. It also satisfies (E2)-(E4). The remaining step is to show that (E1) holds when $t = L$ and $p_L^m > \underline{p}$. For any p such that $E(p) = 0$, arguments using (SCP) developed in the proof of (ii) below establish that a deviation is non-improving: $V(\underline{p}, 0, L) > V(p, 0, L)$. Among p such that $E(p) = 1$, the most attractive deviation is the monopoly price, p_L^m . It is thus sufficient to confirm that $V(\underline{p}, 0, L) > V(p_L^m, 1, L)$. To this end define $p' < \underline{p}$ by $V(p', 0, H) = V(p_L^m, 1, H)$. Using the concavity of Π and thus V in p , as well as (SCP), we see that this deviation is also non-improving: $V(\underline{p}, 0, L) > V(p', 0, L) > V(p_L^m, 1, L)$.

(ii) Let $\{P, E, b\}$ be a separating intuitive equilibrium. First, (E2) and (E3) imply that $E(P(L)) = 0 < 1 = E(P(H))$. Second, $P(H)$ must be equal to p_H^m , since $P(H) \neq p_H^m$ implies

$$V(P(H), 1, H) < V(p_H^m, 1, H) \leq V(p_H^m, E(p_H^m), H)$$

in contradiction to (E1). Third, since the H incumbent can deter entry by choosing $P(L)$, we must have

$$V(P(L), 0, H) \leq V(p_H^m, 1, H)$$

which implies that $P(L) \notin (\underline{p}, \bar{p})$.

Consider the case $p_L^m \geq \underline{p}$. The concavity of Π and hence of V in p implies that $V(\underline{p}, 0, L) > V(p, 0, L)$ for all $p < \underline{p}$ and $V(\bar{p}, 0, L) > V(p, 0, L)$ for all $p > \bar{p}$. The definition of \underline{p} and \bar{p} together with (SCP) imply $V(\underline{p}, 0, L) > V(\bar{p}, 0, L)$. Therefore, it follows that, if $P(L) \notin [\underline{p}, \bar{p})$, then there is $\varepsilon > 0$ such that

$$V(\underline{p} - \varepsilon, 0, L) > V(P(L), 0, L) \quad \text{and} \quad V(\underline{p} - \varepsilon, 0, H) < V(p_H^m, 1, H)$$

But then (E4) implies that $b_L(\underline{p} - \varepsilon) = 1$. Hence $E(\underline{p} - \varepsilon) = 0$ and $V(\underline{p} - \varepsilon, 0, L) > V(P(L), 0, L)$ means that (E1) fails. Therefore, it must be that $P(L) \in [\underline{p}, \bar{p})$, which together with the previous conclusion that $P(L) \notin (\underline{p}, \bar{p})$ gives $P(L) = \underline{p}$.

The corresponding argument for the case $p_L^m < \underline{p}$ is that, if $P(L) \neq p_L^m$, then

$$V(p_L^m, 0, L) > V(P(L), 0, L) \quad \text{and} \quad V(p_L^m, 0, H) < V(p_H^m, 1, H)$$

Thus, by (E4), $b_L(p_L^m) = 1$ and the incumbent's deviation to p_L^m would be profitable, so $P(L) = p_L^m$.

(iii). Consider the case $p_L^m \geq \underline{p}$ and $b_L \geq \tilde{b}_L$. Let $p' \in [\underline{p}, p_L^m]$ and define $\{P, E, b\}$ as follows: $P(L) = P(H) = p'$; $E(p) = 0$ for $p \leq p'$ and $E(p) = 1$ for $p > p'$; $b_L(p) = b_L^0$ for $p \leq p'$ and $b_L(p) = 0$ for $p > p'$. It is a routine matter to verify that $\{P, E, b\}$ satisfy (E1)-(E3). To verify that b satisfies (E4), observe that all $p < p'$ are sure to reduce profit below the equilibrium level for both types, and hence (E4) places no restriction. Next define p'' by $V(p', 0, H) = V(p'', 0, H)$. For $p \in (p', p'']$, $V(p, 0, H) \geq V(p', 0, H)$, and hence $b_L(p) = 0$ satisfies (E4). For $p > p''$, observe that (SCP) implies $V(p'', 0, L) < V(p', 0, L)$ and then the concavity of Π in p implies $p'' > p_L^m$. Hence $V(p, 0, L) < V(p'', 0, L)$ and consequently $V(p, 0, L) < V(p', 0, L)$, so that $b_L(p) = 0$ satisfies (E4). Therefore, $\{P, E, b\}$ is a pooling intuitive equilibrium.

(iv) Let $\{P, E, b\}$ be a pooling intuitive equilibrium. Let p' denote the equilibrium price. First, $E(p')$ must be 0, since otherwise, for at least one $t \in \{L, H\}$, $p_t^m \neq p'$ and an incumbent of this type t could profitably deviate to p_t^m . Equilibrium profits are thus given by $v(L) \equiv V(p', 0, L)$ and $v(H) \equiv V(p', 0, H)$.

Clearly, $p' \geq \underline{p}$, since otherwise the H incumbent will deviate to p_H^m . Suppose then that $p' > p_L^m$. There are two cases to consider. First, if $p' > p_H^m$, define $p'' < p_H^m$ by $V(p'', 0, H) = v(H)$. Then (SCP) implies $V(p'', 0, L) > v(L)$, and so (E2) and (E4) imply $E(p'' - \varepsilon) = 0$ for small $\varepsilon > 0$. But then $V(p'' - \varepsilon, E(p'' - \varepsilon), L) > v(L)$, contradicting (E1). Second, if $p' \in (p_L^m, p_H^m]$, choose a sufficiently small $\varepsilon > 0$ such that $p' - \varepsilon > p_L^m$. Then $V(p' - \varepsilon, 0, H) < v(H)$ and $V(p' - \varepsilon, 0, L) > v(L)$, and so (E4) and (E2) imply $E(p' - \varepsilon) = 0$. Therefore, $V(p'' - \varepsilon, E(p'' - \varepsilon), L) > v(L)$, contradicting (E1) for the L incumbent. The conclusion is that $p' \leq p_L^m$ and hence $p' \in [\underline{p}, p_L^m]$. ■

Proposition 4.1: *There exists a symmetric SPE which maximizes the total payoff. Along its path, $p_i^t = p_j^t = p(a^t)$, where*

$$\begin{aligned} p(L) = p(H) = 1 & \quad \text{for } \delta > \frac{H}{(1+w)H+(1-w)L} \\ p(L) = 1, p(H) = \frac{\delta(1-w)L}{H[1-\delta(1+w)]} & \quad \text{for } \frac{H}{(1+w)H+(1-w)L} \geq \delta \geq \frac{1}{2} \\ p(H) = p(L) = 0 & \quad \text{for } \delta < 1/2 \end{aligned}$$

Proof. First, let us verify that the path described in the claim is consistent with SPE. This path is the outcome of the following firms' strategies: charge p_i in state $i = L, H$, unless there has been a deviation, in which case charge 0. Obviously, these strategies are mutual best responses in any subgame following a deviation. In other subgames, there are only two relevant deviations to consider:

slightly undercutting $p(H)$ in state H and slightly undercutting $p(L)$ in state L . Undercutting $p(H)$ is unprofitable if and only if

$$\{p(H)H + \delta[wp(H)H + (1-w)p(L)L]/(1-\delta)\}/2 \geq p(H)H$$

where the LHS captures the payoff of continuing along the path and the RHS captures the payoff associated with undercutting (a slight undercutting gives the deviant almost twice the equilibrium profit once and zero thereafter). Similarly, undercutting $p(L)$ is unprofitable if and only if,

$$\{p(L)L + \delta[wp(H)H + (1-w)p(L)L]/(1-\delta)\}/2 \geq p(L)L$$

Now, it can be verified that the p_i 's of the proposition satisfy these two conditions in the appropriate ranges.

The following three steps show that this equilibrium maximizes the sum of the firms' payoffs, over the set of all SPE. First, for any SPE, there is a symmetric SPE in which the sum of the payoffs is the same. To see this, take a SPE in which $p_i^t \neq p_j^t$ somewhere on the path and modify it so that everywhere on the path the two prices are equal to $\min\{p_i^t, p_j^t\}$ and so that any deviation is punished by reversion to the zero prices forever. At t such that in the original SPE $p_i^t < p_j^t$, firm j still does not want to undercut, since its continuation value is at least half while its immediate gain is exactly half of the corresponding gains in the original SPE. By symmetry, this applies to i as well. At t such that in the original SPE $p_i^t = p_j^t$, for at least one of the firms, the continuation value is not smaller while the gain from undercutting is the same as in the original SPE, and by symmetry the other firm does not profit from the undercutting either. Second, let V denote the maximal sum of payoffs over the set of all SPE (since the set of SPE payoffs is compact, such maximum exists). Consider a symmetric equilibrium with sum of payoffs V (which exists by the first step). Observe that V must be the sum of payoffs in any subgame on the path that starts at the beginning of any period t before a^t was realized, i.e., after a history of the form $(a^1, p_i^1, p_j^1), \dots, (a^{t-1}, p_i^{t-1}, p_j^{t-1})$. If it were lower for some t , then the strategies in that subgame could be changed to yield V . This would not destroy the equilibrium elsewhere, since it would only make deviations less profitable. But it will raise the sum of payoffs in the entire game, in contradiction to the maximality of V . Now, after any history along the path of this equilibrium that ends with a^t the equilibrium strategies must prescribe the price

$$\varphi(a^t) = \arg \max_p \{pa^t \text{ s.t. } (pa^t + \delta V)/2 \geq pa^t \text{ and } p \leq 1\} \quad (8.1)$$

Otherwise, the equilibrium that prescribes these prices at t and continues according to the considered equilibrium elsewhere would have a higher sum of payoffs. Therefore, $V = [w\varphi(H)H + (1 - w)\varphi(L)L]/(1 - \delta)$. Upon substituting this for V in (8.1), a direct solution of this problem yields $\varphi(x) = p(x)$, $x = L, H$, where $p(x)$ are given in the proposition. ■

Proposition 5.1:(A) *There does not exist a pure-strategy Nash equilibrium.*

(B) *There exists a unique symmetric Nash equilibrium F . It satisfies:*

(i) $\bar{p}(F) = v$;

(ii) $[\underline{p}(F) - c](U/N + I) = [v - c](U/N)$

(iii) $[\underline{p} - c](U/N + (1 - F(p))^{N-1}I) = [v - c](U/N)$ for every $p \in [\underline{p}(F), \bar{p}(F)]$.

Proof. (A) Let k denote the number of firms selecting the lowest price, p , and begin with the possibility that $2 \leq k \leq N$. If $p > c$, then a low-priced firm would deviate from the putative equilibrium with a price just below p , since $[p - c](U/N + I) > [p - c](U/N + I/k)$. On the other hand, if $p = c$, then a low-priced firm could deviate to $p' > p$ and earn greater profit, since $(p' - c)(U/N) > 0$. Consider next the possibility that $k = 1$. Then the low-priced firm could deviate to $p + \varepsilon$, where ε is chosen so that all other firms' prices exceed $p + \varepsilon$, and earn greater profit, since $[p + \varepsilon - c](U/N + I) > [p - c](U/N + I)$.

(B) We begin by showing that any symmetric Nash equilibrium F satisfies (i)-(iii). First, we note that, by the argument of the previous paragraph, $\underline{p}(F) > c$. We next argue that F cannot have a mass point. If p were a mass point of F , then a firm could choose a deviant strategy that is identical to the hypothesized equilibrium strategy, except that it replaces the selection of p with the selection of $p - \varepsilon$, for ε small. The firm then converts all events in which it ties for the lowest price at p with events in which it uniquely offers the lowest price at $p - \varepsilon$. Since ties at p occur with positive probability, and since $p \geq \underline{p}(F) > c$, the firm's expected profit would then increase if ε is small enough.

Suppose now that $\bar{p}(F) < v$. Given that no price is selected with positive probability, ties occur with zero probability. Thus, when a firm chooses $\bar{p}(F)$, with probability one, it sells only to uninformed consumers. For ε small, the firm would increase expected profits by replacing the selection of prices in the set $[\bar{p}(F) - \varepsilon, \bar{p}(F)]$ with the selection of the price v . Thus, $\bar{p}(F) = v$. Similarly, when a firm selects the price $\underline{p}(F)$, with probability one it uniquely offers the lowest price in the market and thus sells to all informed consumers. Since expected profit must be constant throughout the support of F , it follows that $[\underline{p}(F) - c](U/N + I) = [v - c](U/N)$.

We argue next that F is strictly increasing over $(\underline{p}(F), \bar{p}(F))$. Suppose instead that there exists an interval (p_1, p_2) such that $\underline{p}(F) < p_1$, $\bar{p}(F) > p_2$ and $F(p_1) = F(p_2)$. In this case, prices in the interval (p_1, p_2) are selected with zero probability. For ε small, a firm then would do better to replace the selection of prices in the interval $[p_1 - \varepsilon, p_1]$ with the selection of the price $p_2 - \varepsilon$. Since prices in the interval (p_1, p_2) are selected with zero probability, the deviation would generate (approximately) the same distribution over market shares but at a higher price.

It follows that any interval of prices resting within the larger interval $[\underline{p}(F), \bar{p}(F)]$ is played with positive probability. It thus must be that all prices in the interval $[\underline{p}(F), \bar{p}(F)]$ generate the expected profit $[v - c](U/N)$. Now, the probability that a given price p is the lowest price is $[1 - F(p)]^{N-1}$. Thus, we get the iso-profit equation: $[p - c](U/N + (1 - F(p))^{N-1}I) = [v - c](U/N)$ for all $p \in [\underline{p}(F), \bar{p}(F)]$.

Having proved that (i)-(iii) are necessary for a symmetric Nash equilibrium, we now complete the proof by confirming that there exists a unique distribution function satisfying (i)-(iii) and that it is indeed a symmetric Nash equilibrium strategy. Rewrite (iii) as $[1 - F(p)]^{N-1} = (v - p)U/N(p - c)I$ and observe from (iii) that, for $p \in (\underline{p}(F), \bar{p}(F))$, the RHS is between 0 and 1, so that there is a unique solution $F(p) \in (0, 1)$. It follows from (i)-(iii) that $F(\underline{p}(F)) = 0 < 1 = F(\bar{p}(F))$ and $F'(p) > 0$ for $p \in (\underline{p}(F), \bar{p}(F))$, confirming that F is indeed a well-defined distribution. To verify that F is a Nash equilibrium, consider any one firm and suppose that all other $N - 1$ firms adopt the strategy $F(p)$ defined by (i)-(iii). The given firm then earns a constant expected profit for any price in $[\underline{p}(F), \bar{p}(F)]$, and so it cannot improve upon F by altering the distribution over this set. Furthermore, any price below $\underline{p}(F)$ earns a lower expected profit than does the price $\underline{p}(F)$, and prices above $\bar{p}(F) = v$ are infeasible. Given that its rivals use the distribution function F , the firm can do no better than to use F as well. ■

Proposition 5.2: (i) *In the incomplete-information game, there exists a pure-strategy and strict Nash equilibrium, P .*

(ii) *Given a constant $c \in (0, v)$, for any $\varepsilon > 0$, there exists $\delta > 0$ such that, if $|c(t) - c| < \delta$ for all t , then the distribution of prices induced by P , $P^{-1}(x)$, approaches F_c , the mixed-strategy equilibrium of the complete-information game with common per-unit cost c .*

Proof. (i) Let $P : [0, 1] \rightarrow [c(0), v]$ be defined by the following differential

equation and boundary condition:

$$P'(t) = \frac{[P(t) - c(t)][N - 1][1 - t]^{N-2}I}{U/N + [1 - t]^{N-1}I} \quad (8.2)$$

$$P(1) = v. \quad (8.3)$$

Clearly, such a solution P exists and satisfies $P(t) > c(t)$ and $P'(t) > 0$ for all t .

We next show that P is a symmetric Nash equilibrium strategy. Let $\Psi(t, \tilde{t})$ denote the expected profit of a firm of type t that picks price $P(\tilde{t})$ when its rivals employ the strategy P ,

$$\Psi(t, \tilde{t}) = [P(\tilde{t}) - c(t)]\{U/N + [1 - \tilde{t}]^{N-1}I\}$$

Notice that this formula utilizes the strict monotonicity of P by letting $[1 - \tilde{t}]^{N-1}$ describe the probability that $P(\tilde{t})$ is the lowest price. To verify the optimality of $P(t)$ for a type t firm, we only have to check that $P(t)$ is more profitable than other prices in the support of P (the strict monotonicity of P implies that $P(c(0))$ is more profitable than any $p < P(c(0))$, and $P(1) = v$ is more profitable than any $p > v$). The function P thus constitutes a symmetric pure-strategy Nash equilibrium if the following incentive-compatibility condition holds:

$$\Psi(t, t) \geq \Psi(t, \tilde{t}) \quad \text{for all } t, \tilde{t} \in [0, 1] \quad (8.4)$$

Observe that

$$\Psi_2(t, \tilde{t}) = -[P(\tilde{t}) - c(t)][N - 1][1 - \tilde{t}]^{N-2}I + \{U/N + [1 - \tilde{t}]^{N-1}I\}P'(\tilde{t}) \quad (8.5)$$

It therefore follows from (8.2) that

$$\Psi_2(t, t) = 0 \quad \text{for all } t \in [0, 1]$$

Observe next that

$$\begin{aligned} \Psi(t, t) - \Psi(t, \tilde{t}) &= \int_{\tilde{t}}^t \Psi_2(t, x) dx = \int_{\tilde{t}}^t [\Psi_2(t, x) - \Psi_2(x, x)] dx = \int_{\tilde{t}}^t \left(\int_x^t \Psi_{12}(y, x) dy \right) dx \\ &= \int_{\tilde{t}}^t \left(\int_x^t c'(y)(N - 1)[1 - x]^{N-2} I dy \right) dx > 0 \end{aligned}$$

where the second equality follows from $\Psi_2(x, x) = 0$ and the expression for $\Psi_{12}(y, x)$ is obtained by differentiating (8.5). Therefore, (8.4) is satisfied and

this establishes that the pure strategy P defined above gives a Nash equilibrium. Notice further that a firm of type t strictly prefers the price $P(t)$ to any other.

(ii) To establish the approximation result, let $c \in (0, v)$ and let F_c denote the symmetric mixed-strategy equilibrium strategy in the complete-information game with common per-unit costs c . Define the function P_c by

$$P_c(t) = F_c^{-1}(t) \quad \text{for } t \in [0, 1]$$

This definition means that the distribution of prices induced by P_c is the same as the distribution of prices generated by the equilibrium mixed-strategy F_c of the complete-information game. Observe that P_c is the solution to (8.2)-(8.3) for $c(t) \equiv c$. First, note that $P_c(1) = F_c^{-1}(1) = v$. Next, differentiate the identity given in part B(iii) of Proposition 5.1 to get

$$1 = \frac{[p - c][N - 1][1 - F(p)]^{N-2}F'(p)I}{U/N + [1 - F(p)]^{N-1}I} \quad (8.6)$$

Multiply both sides of (8.6) by $P'_c(t)$ and substitute there $p = P_c(t)$, $F = F_c$ and $t = F_c(P_c(t))$ to get

$$P'_c(t) = \frac{[P_c(t) - c][N - 1][1 - t]^{N-2}I}{U/N + [1 - t]^{N-1}I}$$

So the function P_c solves (8.2).

Next observe that (8.2)-(8.3) define a continuous functional, ϕ , from the space of non-decreasing cost functions, $c : [0, 1] \rightarrow (0, v)$, into the space of non-decreasing price functions, $P : [0, 1] \rightarrow [0, v]$. Thus, for an increasing function $c(\cdot)$ the price function $P(\cdot) = \phi(c(\cdot))$ is the symmetric equilibrium in the incomplete-information game with costs $c(\cdot)$, while for $c(\cdot) \equiv c$, $\phi(c) = P_c$. Therefore, invoking the continuity of ϕ , we conclude that, for any $\varepsilon > 0$, there exists $\delta > 0$ such that if $\|c(\cdot) - c\| < \delta$, then $\|\phi(c(\cdot)) - P_c\| < \varepsilon$. In other words, the pure-strategy Nash equilibrium that arises in the incomplete-information game generates approximately the same distribution over prices as occurs in the mixed-strategy equilibrium of the complete-information game. ■

References

- [1] Abreu, D., D. Pearce and E. Stacchetti (1985), "Optimal Cartel Equilibria with Imperfect Monitoring," *Journal of Economic Theory*, 39, 251-269.
- [2] Areeda, P. and D. Turner, "Predatory Pricing and Related Practices Under Section 2 of the Sherman Act," *Harvard Law Review*, 88 (December 1975), 697-733.
- [3] Bagwell, K. and G. Ramey (1988), "Advertising and Limit Pricing," *Rand Journal of Economics*, 19, 59-71.
- [4] Bagwell, K. and G. Ramey (1991), "Oligopoly Limit Pricing," *Rand Journal of Economics*, 22, 155-72.
- [5] Bagwell, K and R. Staiger (1997), "Collusion over the Business Cycle," *Rand Journal of Economics*, 28, 82-106.
- [6] Bain, J. (1949), "A Note on Pricing in Monopoly and Oligopoly," *American Economic Review*, 39: 448-64.
- [7] Borenstein, S. and A. Shephard (1996), "Dynamic Pricing in Retail Gasoline Markets," *Rand Journal of Economics*, 27, 429-451.
- [8] Brander J. and B. Spencer (1985), "Export Subsidies and International Market Share Rivalry," *Journal of International Economics*, 18, 83-100.
- [9] Bulow, J., J. Geanakoplos and P. Klemperer (1985), "Multimarket Oligopoly: Strategic Substitutes and Complements," *Journal of Political Economy*, 93 (June 1985), 488-511.
- [10] Cho, I-K. and D. Kreps (1987), "Signalling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102: 179-221.
- [11] Dixit, A. (1980), "The Role of Investment in Entry Deterrence," *Economic Journal*, 90, 95-106.
- [12] Eaton, J. and G. Grossman (1986), "Optimal Trade and Industrial Policy Under Oligopoly," *Quarterly Journal of Economics*, 101, 383-406.

- [13] Fellner, W. (1949), *Competition Among the Few*, New York: Knopf.
- [14] Friedman, J. (1990), *Game Theory with Applications to Economics*, Oxford University Press.
- [15] Friedman, J. (1971), "A Non-cooperative Equilibrium in Supergames," *Review of Economic Studies*, 38.
- [16] Fudenberg, D. and J. Tirole (1984), "The Fat Cat Effect, the Puppy Dog Ploy and the Lean and Hungry Look," *American Economic Review, Papers and proceedings*, 74, 361-368.
- [17] Fudenberg, D. and J. Tirole (1986), *Dynamic Models of Oligopoly*. Chur: Harwood Academic Publishers.
- [18] Fudenberg, D. and J. Tirole (1987), "Understanding Rent Dissipation: On the Use of Game Theory in Industrial Organization," *American Economic Review, Papers and proceedings*, 77, 176-183.
- [19] Fudenberg, D., D. Levine and E. Maskin (1994), "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62, 997-1039.
- [20] Green, E. and R. Porter, "Noncooperative Collusion Under Imperfect Price Information," *Econometrica*, 52 (January 1984), 87-100.
- [21] Haltiwanger, J. and J. Harrington (1991), "The Impact of Cyclical Demand Movements on Collusive Behavior," *Rand Journal of Economics*, 22, 89-106.
- [22] Harrington, J. (1986), "Limit Pricing when the Potential Entrant is Uncertain of Its Cost Function," *Econometrica*, 54: 429-37.
- [23] Harsanyi, J. (1967-68), "Games with Incomplete Information Played by 'Bayesian' Players," Parts I, II, and III, *Management Science*, 14, 159-82, 320-24, 486-502.
- [24] Harsanyi, J. (1973), "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points," *International Journal of Game Theory*, 2: 1- 23.
- [25] Kadiyali, V., "Entry, Its Deterrence, and Its Accommodation: A Study of the U.S. Photographic Film Industry," *Rand Journal of Economics*, 27.3 (Autumn 1996), 452-478.

- [26] Kreps, D. and J. Scheinkman, “Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes,” *Bell Journal of Economics*, 14 (Autumn 1983), 326-337.
- [27] Kreps, D. and R. Wilson (1982a), “Sequential Equilibria,” *Econometrica*, 50: 863- 894.
- [28] Kreps, D. and R. Wilson (1982b), “Reputation and Incomplete Information,” *Journal of Economic Theory*, 27: 253-79.
- [29] McGee, J. (1958), “Predatory Price Cutting: The Standard Oil (N.J.) Case,” *Journal of Law and Economics*, 1, 137-169.
- [30] Milgrom, P. and J. Roberts (1982a), “Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis,” *Econometrica*, 50: 443-59.
- [31] Milgrom, P. and J. Roberts (1982b), “Predation, Reputation and Entry Deterrence,” *Journal of Economic Theory*, 27: 280-312.
- [32] Modigliani, F. (1958), “New Developments on the Oligopoly Front,” *Journal of Political Economy*, 66, 215-232.
- [33] Nash, J. (1950), “Equilibrium Points in n-person Games,” *Proceedings of the National Academy of Sciences*, 36: 48-49.
- [34] Porter, R. (1983), ”A Study of Cartel Stability: The Joint Executive Committee, 1880-1886” *Bell Journal of Economics*, 14, 301-314.
- [35] Radner, R. (1981), “Monitoring Cooperative Agreements in a Repeated Principal–Agent Relationship,” *Econometrica*, vol 49, No. 5, 1127-1148.
- [36] Roberts, J. (1985), “A Signaling Model of Predatory Pricing,” *Oxford Economic Papers, Supplement*, 38, 75-93.
- [37] Rosenthal, R. (1980), “A Model in Which an Increase in the Number of Sellers Leads to a Higher Price,” *Econometrica*, 48, 1575-1580.
- [38] Rotemberg J. and G. Saloner, “A Supergame-Theoretic Model of Business Cycles and Price Wars During Booms,” *American Economic Review*, 76 (June 1986), 390-407.

- [39] Rubinstein, A. (1979), “Offenses that May Have Been Committed by Accident—An Optimal Policy of Retribution,” in S. Brams, A. Schotter and G. Schwodiauer (eds), *Applied Game Theory*: 236-53. Wurzburg, Vienna: Physica-Verlag.
- [40] Selten, R. (1965), “Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragertragheit,” *Zeitschrift für die gesamte Staatswissenschaft*, 12, 301–324.
- [41] Shilony, Y. (1977), “Mixed Pricing in Oligopoly,” *Journal of Economic Theory*, 14: 373-88.
- [42] Spence, M. (1977), “Entry, Capacity, Investment and Oligopolistic Pricing,” *Bell Journal of Economics*, 8, 534-544.
- [43] Stigler, G. (1964), “A Theory of Oligopoly,” *Journal of Political Economy*, 72, 44-61.
- [44] Tirole, J. (1988), *The Theory of Industrial Organization*, Cambridge: MIT Press.
- [45] Varian, H. (1980), “A Model of Sales,” *American Economic Review*, 70, 651-659.
- [46] Villas-Boas, J. Miguel (1995), “Models of Competitive Price Promotions: Some Empirical Evidence from the Coffee and Saltine Crackers Markets,” *Journal of Economics and Management Strategy*, 4: 85-107.