10. The Economics of Industry Group at STICERD

John Sutton

No field of economics has changed more dramatically over the past twenty years than has Industrial Organisation ("Economics of Industry"). When STICERD was founded, this field was dominated by purely empirical studies, which rested on a rather sketchy theoretical base. The subject matter of the field was concerned with asking how the different patterns of technology and tastes found in different markets affect the modes of competition in these markets, and the way in which they developed over time. How, for example, do firms compete in terms of their advertising and marketing activities? What determines their levels of R&D expenditure? Why are some industries dominated globally by a handful of firms? How does price competition work in such concentrated industries? The traditional Industrial Organisation literature tackled such questions, for the most part, by designing cross-industry studies in which differences in outcomes could be associated with measurable 'industry characteristics'. This literature uncovered a substantial body of apparent 'statistical regularities', but the theoretical underpinnings of the work were often weak.

1. A New Literature

The foundation of STICERD coincided with the upsurge of a new wave of interest in such questions, which arose from developments within applied game theory. From the end of the 1970s, a new theoretical literature developed which held centre-stage within the Industrial Organisation field for a decade. This new literature brought together theorists and Industrial Organisation specialists, and much of the new research effort was driven by attempts to specify formal models of phenomena which had long been familiar at the empirical level, but which were poorly understood analytically. During the 1980s, the theory group at STICERD had been very active in this literature. Dasgupta's 1980 paper with Stiglitz on process innovation opened up a new way of looking at the determinants of market structure. The sequence of papers by Shaked and Sutton (1983, 1984a,b, 1987), together with the parallel work of Jean Gabszewicz and Jacques Thisse at CORE, developed the analysis of vertical product differentiation, which would later provide a basis for the study of product innovation. The area began to attract increasing numbers of graduate students in the early 1980s, and notable theses from the period include Giacomo Bonnano's work on entry and pre-emption, and the work of Tilman Borges on multistage games with sequential entry. By the mid-80s, the new field was thriving within STICERD - yet it was already becoming evident that, at the very core of the enterprise, a rather worrying question was emerging.

The question was prompted, paradoxically, by the very success which had attracted so many researchers to the field in the early 1980s. The new stage-game models could be designed in many alternative ways, and often there was no *a priori* reason to prefer one design over another. Thus, for example, in modelling the entry of firms to a market, it might be assumed that the firms entered 'simultaneously', or that they entered 'sequentially'. Behind this innocuous-looking distinction lay some quite subtle issues. If, for example, firms A and B entered sequentially, this implied that firm B was allowed to employ a strategy in which its decisions on, say, how much production capacity to build, could be made conditional on what A had done in the preceding stage. Firm A, on the other hand, anticipating B's reaction, would be in a position to manipulate B's choice of capacity level. Now this kind of asymmetry in the positions of A and B would be eliminated if we were to model the situation using a 'simultaneous entry' model: here, each firm would be assumed to take as given his rival's level of capacity, and to choose his own level accordingly.

This 'entry' example, simple though it is, is already sufficient to expose the nature of the central difficulty that arose in the new game-theoretic literature:

the results of the analysis vary dramatically with the form of the entry process that we assume, but it may be very difficult to defend either choice over the other on *a priori* grounds.

The entry process is just one example of this kind of difficulty. Another area where such problems abound is that of modelling price competition. Here, a number of classical models were available from the early literature on oligopoly, the most widely known being the Bertrand and Cournot models. Here again, the choice of either of these models - or of any other model of price competition - was very difficult to justify.

As researchers became aware of these problems during the late 1980s, two lines of response emerged at the empirical level. The first response lay in moving away from the traditional agenda of 'cross-industry studies' and focussing instead on the modelling of a single industry. The motivation for this approach lies in the argument that, with a deep knowledge of some particular market, we may be able to partially resolve the difficulty of model choice by using known facts about the industry to eliminate some models as 'unreasonable', and so narrow the range of choice. Then, we can proceed to use a loosely specified model that incorporates a number of free parameters, and 'let the data decide' what values to place on those parameters. In other words, we travel as far as we can on the basis of *a priori* appeals to what kind of model is reasonable, and after that we engage in a 'model selection' exercise in which some 'goodness-of-fit' criteria are used to argue for some 'preferred specification'.

During the late 1980s, single industry studies designed in this way became increasingly popular. The first PhD thesis produced within the newly formed Economics of Industry group looked at competition within the European car market, and examined the way in which increasing competition from Japanese producers altered the pattern of prices and the structure of the market shares across different countries (Mertens (1990)). The new game-theoretic models provided a rich framework within which to conduct this kind of investigation, and this type of study has continued to play a major role in the Industrial Organisation literature during the past decade. Yet this line of attack has two shortcomings. First, it fails to address the underlying anxiety surrounding the use of game-theoretic models: does the fact that these models can be designed in many ways, leading to a wide range of alternative implications, simply imply that the whole enterprise is empirically empty? In other words, is it true that in these models, 'anything can happen'? If this is the case, then the whole enterprise is indeed empty.

The second objection to the new single industry studies is that this trend implies an abandonment of the traditional agenda of the subject. Traditionally, the field of Industrial Organisation was concerned with the investigation of statistical regularities that appeared in cross-industry studies. Regularities of this kind are interesting for two reasons. First, in allowing us to relate market outcomes to measurable characteristics of different markets, they may help us to uncovering the workings of some basic market mechanisms. Second, the fact that such regularities appear in datasets that span different industries, between which there are endless detailed differences, suggests that these regularities are driven by mechanisms that are sufficiently robust to override many secondary features that differ widely from one industry to the next. Put more generally, the strength of the traditional agenda lay in its insistence on the search for valid generalisations: while accepting that every industry has some features that are quite idiosyncratic, the tradition remained focussed on the question of whether it was possible to make general statements about the general run of industries, or about some broad classes of industries that shared some basic characteristics. Indeed, unless some claims can be made at this level, it is difficult to see how the findings of Industrial Organisation economists can feed back into other areas of economics.

2. The Bounds Approach

These considerations motivate an alternative line of attack in moving from the new generation of game-theoretic models towards an empirical agenda. This new line of attack, introduced in Sutton (1991), builds on the idea that we should not think in terms of specifying or identifying some 'true model' of each market, but rather we should aim only to identify some broad class of 'admissible models' i.e. the set of models that cannot be ruled out a priori as a representation of the market (an idea introduced in a 1984 STICERD paper, published as Shaked and

Sutton (1987)). To see how this theoretical idea relates to the empirical agenda, it is useful to look back at the emerging problems in the Industrial Organisation area from the perspective of an econometrician.

To the econometrician, the fact that theory offers only a broad menu of candidate models is not in itself a problem. In principle, all that is required is that all the candidate models should be embedded in a 'grand model' that contains many parameters: switching these parameters on or off brings us to different candidate models. Insofar as we wish to propose that outcomes will depend on a list of 'industry characteristics', all that is required in principle is that all such characteristics should be built into the 'grand model', and that we should, in carrying out our empirical investigations, take care to control for the influence of all these industry characteristics.

The problem with game-theoretic models is this: the theoretical literature tells us that some of the industry characteristics that will exert large and systematic influences on outcomes are of a kind that can not be measured, proxied or controlled for in the context of cross-industry studies. The pattern of entry is one such characteristic - we might, for some particular industry, justify the use of some special representation of the entry pattern, but if we are examining statistical regularities that hold across some large group of industries, we cannot in practice hope to find some measurable parameter that serves to control for the fact that a market had one or another kind of 'strategic assymmetry' between early entrants.

In price competition, similar problems arise: no measurable industry characteristics tell us whether a Bertrand model or a Cournot model, or some one of the many alternative models, is appropriate to each market. Yet these features of the markets, which we must of necessity treat as 'unobservables' for the purpose of cross-industry studies, will exert large and systematic effects on outcomes.

So how can we proceed? In the 'bounds' approach of Sutton (1991,1998) the idea is to begin by moving away from the standard paradigm in which the market is represented by a fully-specified model that defines a unique equilibrium outcome (Figure 1a). Instead, we begin from the notion of a 'class of models'. The idea is that models within this class differ from one another in ways that we cannot hope to pin down by reference to measurable market characteristics. All models in the class share certain features, however, and the aim of the theory is to develop propositions that must hold good for any model that has these features, irrespective of those secondary factors that differ between one model and another within the specified class. The aim of all this is to allow us to employ a looser framework than that of the standard paradigm shown in Figure 1a; rather than pin down a unique outcome, the aim is to partition the space of outcomes into those that can be supported as a (perfect Nash) equilibrium of some model within the specified class of models, and those outcomes that cannot be supported in this way (Figure 1b).

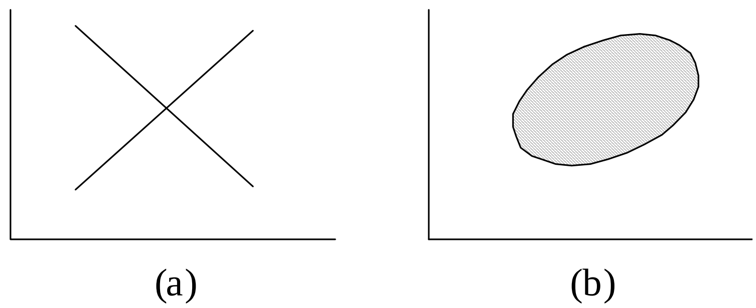

Figure 1. (a) The Standard Paradigm: a unique equilibrium outcome. (b) The Bounds Approach: the shaded area shows a set of outcomes that may be supported as an equilibrium of some 'admissible model'.

In Sutton (1991, 1998), this approach is developed within the context of 'explaining market structure'. The class of models used throughout most of the analysis is of the familiar 'stage game' type: here, firms make a series of investments (in plant and equipment, in advertising or in R&D) over a series of stages. Once these investments have been made, the outcome takes the form of a description of the firm's 'capabilities', which - depending on the context - might be their levels of production capacity, the strength of their 'brand images' as represented by consumers' willingness-to-pay, or their different levels of technical competence resulting from R&D expenditures. What matters is simply that some such kind of summary description of capabilities can be written down. Once this is done, we can suppose that the firms compete with each other - their capabilities being taken as given - and we can write down a payoff function for each firm that specifies its profit in the ensuing ('price competition') subgame as a function of its capability and the capabilities of its rivals.

From a technical point of view, the trick is to move away from the specification of (perfect Nash) equilibria within a particular fully-specified model, towards an equilibrium concept that is defined directly on the space of outcomes. The analysis begins with the introduction of such an equilibrium concept: this rests on two principles. The first is a form of survivor principle: it says that no firm chooses a set of actions (investments) that yield a strictly negative payoff (the 'Viability Condition'). The second principle says that, given the set of actions (investments) made by all 'active' firms (i.e. those that have made some investment), then there does not remain any 'gap in the market', i.e. it would not be possible for some one of the hitherto inactive firms to choose some set of actions (investments) that would yield it strictly positive profits (the 'Stability Condition').

It turns out that the application of these two simple principles suffices to define a set in the space of outcomes (the 'set of equilibrium configurations'), and it further turns out - if we introduce some very mild regularity conditions that exclude certain pathological cases - that this set includes all outcomes that can be reached as (perfect Nash) equilibria of any model within the class we have defined. (In other words, if we choose to model entry as 'simultaneous' we will get one outcome, while if we model it as 'sequential' we will get a different outcome. Both these outcomes will, however, fall within the set of equilibrium configurations, as will any outcomes that can be arrived at by designing the entry process in any other way.

So where does this bring us to in modelling market structure? The main result that emerges is a simple theorem that introduces a new way of looking at the determinants of concentration - and in so doing, provides a simple unified treatment of several hitherto unrelated themes in the literature.

3. Explaining Concentration

While the literature has cited a host of mechanisms that might lead to high concentration (increasing returns to R&D, learning by doing, network effects, etc.), what a game-theoretic approach suggests is that we should tackle a question of this kind using a particular form of argument which is both simple and powerful. If we want to ask, 'why is industry X concentrated?', we begin by turning the question: if industry X was fragmented, in the sense that it contained a large number of firms, each with a small share of industry sales, then would this configuration admit of a profitable deviation by some firm?

Suppose, for example, we consider a market in which firms can raise consumers' willingness-to-pay for their products by spending on R&D. If there are many firms, each with a small market share, then for any given size of market (as measured by the size of the population of consumers, S, say) the amount each firm can afford to spend on R&D while satisfying the viability condition is limited. Now suppose a 'deviant' firm, hitherto inactive, chooses to outspend all the hitherto active firms by spending K times as much on R&D as the highest-spending of these firms. How great a profit will it achieve? Let us denote the profit it will attain in this way by $S\pi(K;\bullet)$, where S denotes the size of the market, as measured by the number of consumers[1]. Now the payoff $S\pi(K;\bullet)$ will depend not only on K, but also on the configuration of investments that have been made by rival firms (denoted by the dot in the function). The idea is to ask: is there anything we can say about $\pi(K;\bullet)$, without knowing anything about this pattern of investments? Let us suppose, for example, that $\pi(K;\bullet)$ is bounded below by some number a(K), so that whatever the spending pattern of rival firms, our deviant firm achieves a gross payoff $S\pi(K;\bullet) \geq aS$ once its spending on R&D exceeds that of any rival firm by a factor of K. The main theorem then tells us that, if the market is sufficiently fragmented, this deviation will prove profitable. It follows that there is some limit to how fragmented the market can become: too low a level of concentration will lead to a violation of the twin conditions of 'viability' and 'stability'. Specifically, it turns out that to avoid such a violation there must be at least one firm in the market that enjoys a share of total industry sales revenue in excess of the ratio a(K)/K. This, in other words, sets a lower

[1] It is customary in this literature to assume flat marginal cost curves, and to suppose that increases in the size of the market come about by way of a replication in the population of consumers, so that the distribution of tastes remains unchanged it then follows that the profit function takes this form.

bound to the (1-firm sales) concentration ratio[2].

From this opening result, the theory can be elaborated in various directions. The first step is to get around the fact that the numbers a and K would be practice by difficult to measure. To by-pass this difficulty, it is necessary to reformulate the result in a different way. The intuition is simple: if R&D is 'ineffective' in raising consumers' willingness-to-pay, little will be spent on R&D. If R&D is effective, however, deviations by 'high spenders' will become profitable. This leads to a more empirically useful statement: if we compare two sets of industries, in one of which R&D spending is (close to) zero, and in the other of which R&D spending exceeds a certain positive fraction of industry sales revenue, then the level of concentration may take values arbitrarily close to zero in the first group, while for the second group the level of concentration must lie above some strictly positive bound.

These simple predictions lend themselves to easily implemented tests. One form of test, used in Sutton (1991) in relation to advertising-intensive industries (where the appropriate theoretical model is similar to that used for R&D competition) involves a comparison between two sets of industries, one group in which the advertising-sales ratio exceeds a certain threshold (1%; in practice, most are >>1%) and a reference group, where the ratio is very small (1%; in practice, very much less than 1%). The test focuses on examining the relation between market size and the level of concentration. The two groups are taken from the Food and Drink sector, and from the six largest Western economies. For the reference group, in which advertising (and R&D) play a negligible role, the viability condition imposes a simple lower bound on the level of concentration, which declines monotonically to zero as market size increases[3] (Figure 2a).

---

[2] In fact, since we can apply this argument for any value K, the bound actually takes the form $C_1 \geq \sup_K a(K)/K$.

[3] To control for the fact that different industries involve different levels of setup cost, market size is measured as a ratio S/ε, where ε is a measure of the cost of a single production plant of 'minimum efficient scale'.

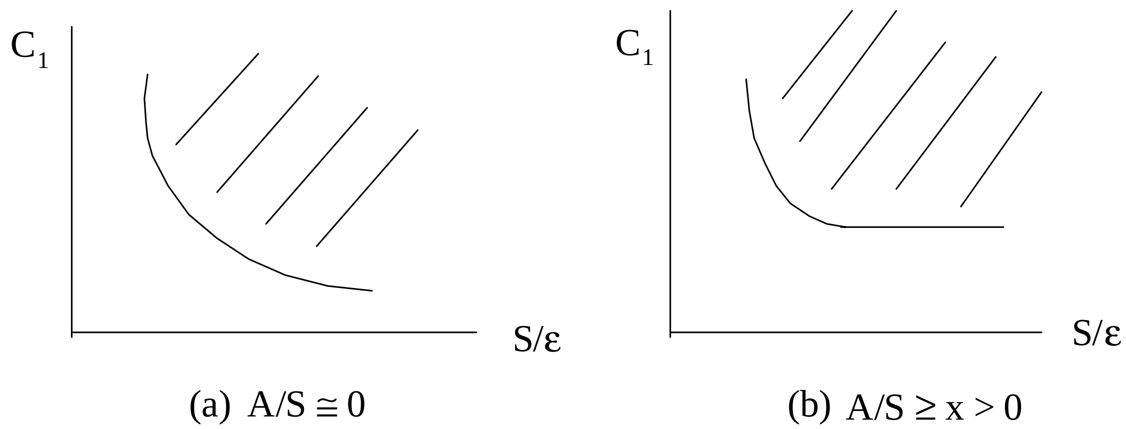

Figure 2. Concentration versus market size.

For the advertising intensive group, on the other hand, the 'nonconvergence' result outlined above implies that there is a lower bound to concentration, which is independent of market size. Here, as market size increases, the effect is not to attract new entrants to the industry, but to spark off increased spending on fixed (advertising) outlays by active firms (Figure 2b).

Empirical tests of this proposition were presented in Sutton (1991), and were strongly supportive of the proposition. Later attempts to replicate the tests on new datasets for the U.S. (Robinson and Chiang (1996)) and for Europe (Lyons and Matraves (1996)) provided confirmation of this finding.

Why is this simple relationship important? There are three reasons:

(i) Encompassing: Traditional investigations of the determinants of market structure were based on regressions of concentration on such measures as S/ε and A/S. The finding was that C1 was negatively related to S/ε, and that there was a very weak positive link to A/S. Under the present theory, such a regression is a misspecification, but it can be shown that, if this theory is correct, such a misspecified regression will indeed yield the traditional results (Sutton (1991) pp.127-8).

(ii) Unifying: While I have so far described the basic nonconvergence theorem in a context where firms spend R&D or advertising outlays to raise consumers' willingness-to-pay for their products, the basic argument can be rephrased to fit such models as 'learning-by-doing' or 'network effects' models. In these models, a firm sets a low price in some 'early' stage in order to attain a high level of output, or sales. The higher the first period level of output or sales, the lower is the firm's level of cost in some later period ('learning effects'), or the higher is consumers' willingness-to-pay for its product ('network effects'). In this setting, the same nonconvergence result operates once again: all we have to do is to replace references to the firm raising its fixed outlays in stage 1, with references to the opportunity cost of profits foregone in stage 1, as the firm builds up its output levels (or customer base) in that period (Sutton (1998), Chapters 14,15). In this way, the basic mechanism leading to high concentration is seen to be the same, across a wide range of contexts hitherto thought of as basically different.

(iii) Extensions: Once this simply nonconvergence result is established, it is possible to extend it to more complex settings, and thereby tackle a range of classical questions regarding market structure. Some of these extensions are described in the next three sections.

4. Price Competition

Up to this point, I have focused attention on the nonconvergence theorem, while leaving in the background the relationship between market size and concentration in those industries where advertising and R&D play no substantial role (see Panel (a) of Figure 2 above). Here, there is a lower bound to concentration which falls to zero as market size goes to infinity ('convergence'). A central theorem within the bounds approach links the height of this lower bound to the 'toughness of price competition' in the industry.

An increase in the toughness of price competition[4] , arising for example from the introduction of legal restrictions on cartels or other forms of price coordina-

[4]The phrase 'toughness of price competition' relates, not to the level of prices or unit margins, which is an endogenous outcome in these models, but rather to the functional relationship between concentration and the level of prices or unit margins. The phrase summarises the nature of competition in the 'final stage subgame', where firms - taking competences as given, compete in price. Thus, in terms of the elementary models, a Cournot model would correspond to 'less tough' price competition than a Bertrand model. At the empirical level, exogenous changes such as the introduction of anti-cartel laws, or the building of transport links between hitherto separated local markets, will alter the toughness of price competition in the market (Sutton (1991), Chapter 6).

tion, is predicted to lead to an upward shift in the lower bound to concentration. The intuition here is straightforward: the viability condition requires that firms' gross profits cover their fixed outlays. If the toughness of price competition increases, this would - for any given level of concentration - reduce firms' profit. Only by a compensating rise in concentration can viability be restored[5] .

Some 'natural experiments' relating to this shift were presented in Sutton (1991), Chapter 6. A particularly noteworthy example of such a natural experiment arose in the UK during the 1960s, where two rounds of legislative changes severely constrained the use of devices to coordinate prices that had hitherto been widely used. In a recent thesis, and a series of STICERD papers, Symeonides (1997a,b) has shown how these legislative changes led to an upward shift in the level of concentration in UK manufacturing industry[6] .

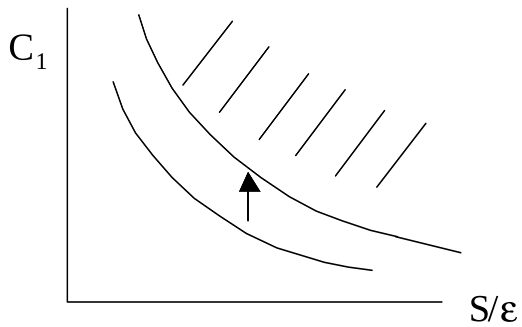

Figure 3. An increase in the toughness of price competition.

5. Dynamic Models

A second strand of recent research in the industry group relates to the extension of the 'stage-game' framework, used throughout most of the literature, to the more complex setting of dynamic games. In this setting, firms compete over successive dates t = 1, 2, 3 and at each date, they choose some level of fixed

[5]In the short run, this effect may be offset insofar as some of the fixed outlays represented by e are 'sunk'. Only in the long run, as plant and equipment needs to be replaced, will the viability condition 'bite'. See Sutton (1991), Chapter 6.

[6]Figure 3 relates to industries where advertising and R&D are unimportant. Where such fixed outlays play an important role, the theoretical link between market size and concentration is slightly more subtle (see Symeonides (1997a and b) for details).

outlays that raises their level of competence; profits accrue as a flow over time, the profits in any period being determined by firms' competence in that period.

If we remain within a simple deterministic setting, in which a given (increment to) fixed outlays determines a corresponding (increment to) a firm's competence, then the same kind of nonconvergence theorem can be formulated, albeit at the cost of a considerable increase in technical complexity (see for example, Sutton (1998), Chapter 13 and the recent Economics of Industry discussion paper of Nocke (1998)).

Matters become more complex when we move to a stochastic setting, in which the level of fixed outlays in any period determines the probability with which a firm jumps to a higher level of competence (as measured by consumers' willingness-to-pay). A simple but powerful framework for analysing this kind of setup has been developed by Ariel Pakes and his several co-authors, whose algorithm allows the explicit computation of perfect Nash equilibria (Ericsson and Pakes (1995); Pakes, Gowrisankaran, and McGuire (1993)). In a recent thesis at the Economics of Industry group, Alison Oldale (1997) has used these methods to re-examine the nature of the nonconvergence phenomenon in a setting where the link from R&D spending to enhanced competence is described by a stochastic version of the 'Cournot' model of Sutton (1991). This simple model involves a parameter which measures the 'effectiveness' of R&D spending, and which maps into a lower bound to concentration. In Oldale's analysis, this bound is no longer deterministic. Instead, we find that, given the value of this parameter, there is a minimal level of concentration such that, if we fall below this level, we tend to bounce back 'quickly'. Obtaining a clear analytical characterisation of this phenomenon is difficult, and is the subject of continuing research. Nonetheless, it seems clear from Oldale's analysis that a nice theoretical analogue of the standard nonconvergence theorem is probably available in this setting.

6. Markets and Submarkets

Another recent line of work relates to extending the analysis beyond the setting in which the market is a single well-defined market of the classical kind (all products in the market are substitutes, and all products outside the market are not). In practice, markets are seldom so well-defined as this. Particularly in R&D intensive industries, the typical market will encompass various submarkets. Within each sub-market, products are close substitutes. The various sub-markets may be linked more or less strongly, both on the demand side (via substitutability) and on the supply side (via scope economies in R&D). By developing the theory in this more complex setting, more powerful predictions become possible as to the lower bound to concentration in R&D-intensive industries.

Once this extension to 'many submarkets' is accomplished, it is natural to ask: what happens if we have a market that contains many submarkets, between which the links are weak? This raises a rather fundamental theoretical question, for the standard approach in game-theoretic Industrial Organisation has been to focus on the analysis of 'strategic interactions', and so on markets where any action by any firm affects the payoff of all firms. In this new setting, however, the focus lies in analysing situations where 'independence effects' become important. This raises the question: how can we, in the context of a game-theoretic analysis, incorporate the idea that actions taken in one submarket do not impinge on the profits earned by firms in another submarket. It is argued in Sutton (1998) that this context demands the introduction of a third principle (in addition to the two introduced earlier). What this principle does, is to impose a form on symmetry requirement, in the manner of Harsanyi and Selten (1988). It turns out that the introduction of this principle leads to a series of new predictions, which relate to the way the size distribution of firms evolves over time in markets that contain many (approximately) independent submarkets (Sutton (1998) Chapters 9-12). These predictions, which were tested for U.S. and German data in Sutton (1998), have been further investigated in contexts are varied as Spanish savings banks (de Juan (1999)), retail grocery markets (Walsh and Whelan (1998)), and - in a recent thesis within the industry group - the U.S. chemical industry (Koopmans (1995)) .

7. The Future of Industrial Organisation

The developments that I have been discussing up to now grew out of a deep-seated concern with the question of whether the new generation of game-theoretical models in Industrial Organisation might be 'empirically empty'. Ten years later, it has become clear that these models are in fact very rich empirically, though taking an board their implications does require a fundamental re-assessment of the ways in which we can move from theory to testing in this field. Another thing which has gradually become clear over the past decade is the deep complementarity that holds between the three main types of empirical research that have run along in parallel during the 1990s. The best way of seeing how the field now stands is by thinking of three lines of attack emerging as responses to the different levels of difficulty that we face in looking at different empirical contexts.

The 'single-industry study' approach works best in those special settings in which special institutional features of the market allow us to pin down fairly narrowly the kind of model that offers a valid representation of the market. The most successful example of this approach lies in the area of auctions: here, the rules of the auction specify exactly the design of the game we use, and the problem of 'model selection' disappears. In settings less special than that of auctions the idea is that, while we cannot hope to pin down a single candidate model, we may at least hope to narrow the class of candidate models to the point where we can narrow our predictions to some small set of 'possible outcomes'.

At the other end of the spectrum from this approach lies the 'bounds approach to market structure'. Here, the aim is to widen the domain of application to the general run of markets, knowing that this will entail accepting a quite wide 'class of models', as being reasonable *a priori*. We aim to place only fairly weak constraints on outcomes, while achieving a high level of generality in that the results can be applied across the general run of markets. Clearly, this is complementary to the 'single industry study' approach: for once we have a set of weak results that apply to all industries, we can now add further assumptions to the theory that are valid for specific types of industry, or even for single industries, in an attempt to sharpen the theoretical predictions, at the cost of narrowing the domain of application. (For an example of this strategy, see Sutton (1991), Chapter 9).

Now this discussion might suggest a rather happy state of affairs, in which all studies lay in some 'possibility frontier' in the space of 'Precision of Predictions vs. Breadth of applicability'(see Figure 4 and Sutton (1998)). There is, however, a third possibility, and this brings us to a third, and crucial, strand in the recent literature. One of the lessons of game-theoretic Industrial Organisation is that there are certain settings in which the competitive mechanism simply fails to pin down a single 'equilibrium outcome'. No matter how we add institutional detail to our models, we will never be able, by reference to 'observable industry characteristics', to pin down a unique equilibrium outcome. Rather, the models we see as appropriate admit of many possible equilibria. The 'folk Theorem' literature is the most dramatic example of this problem (Tirole 1990). A related problem arises in the field of 'cartel stability'. Here, we know from the literature that there is an endless variety of possible mechanisms that may lead to the emergence of fairly stable prices, above the 'competitive' level. Generalising about such arrangements is an unrewarding exercise. The number of factors that affect stability is extremely large[7] .

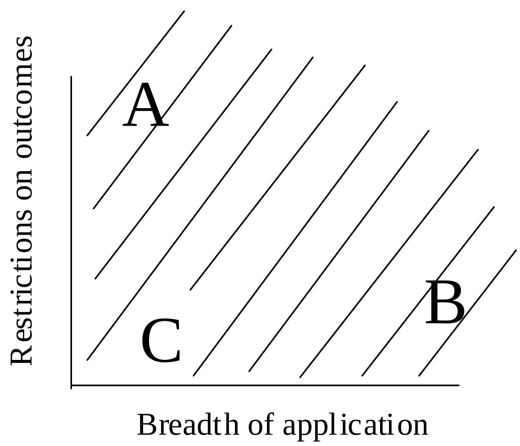

Figure 4. An ABC of Industrial Organisation (A = Auctions, B = Bounds, C = Cartels).

Even if we narrow our domain of reference to a particular cartel at a particular period, it is often very difficult to decide whether a particular model is a 'good' representation of the cartel in question. It still remains possible that

[7]Some recent theses within the group have extended the literature in this area, by looking at the roles played by product differentiation (Raith (1996)) and by capacity constraints (Nocke (1999)).

some useful general relationships may emerge, between the institutional setting of the industry, and the appearance of specific cartel mechanisms. Nonetheless, much of the interesting research in this area at present lies in looking directly at specific cartels, with a view to uncovering the (often highly specific) mechanisms used by that cartel to maintain prices[8] . All in all, then, the old taunt that 'with oligopoly, anything can happen', is not quite dead. Yet it is not quite true, either. The encouraging theme in the recent Industrial Organisation literature is that, while we cannot hope to pin things down in the manner of the standard paradigm (Figure 1a above), we can make very substantial progress across a wide range of phenomena by adopting a more eclectic approach to modelling methods.

8. References

De Juan, R., (1999), "The Independent Submarkets Model: An Application to the Spanish Retail Banking Market," Working paper, Fundación Empresa Pública.

Ericsson, R. and Pakes, A., (1995), "Markov-perfect Industry Dynamics: A Framework for Empirical work," *Review of Economic Studies*, vol. 62(1): 53-82.

Gruber, H., (1990), "Product Innovation and Persistence of Leadership: Theory with Evidence from the Semiconductor Industry," Economics of Industry Discussion Paper no. EI/2, STICERD, London School of Economics.

Harsanyi, J.C. and Selten, R., (1988), *A General Theory of Equilibrium Selection in Games*, Cambridge, MA: MIT Press.

Koopmans, R., (1995), *Asymmetric Industry Structures: Multiple Technologies, Firm Dynamics and Profitability*, PhD Thesis, London School of Economics.

Levenstein, M., (1993), "Price Wars and the Stability of Collusion: A Study of the Pre-World War I Bromine Industry," NBER Working Paper no. 50, Cambridge, MA.

Lyons, B. and Matraves, C., (1996), "Industrial Concentration," in S. Davis and and B. Lyons, *Industrial Concentration in the European Union: Structure,*

[8] For some very nice recent analyses of such problems see Levenstein (1993) and Scott-Morton (1997).

Strategy and the Competitive mechanism, Oxford: Clarendon Press.

Mertens, Y., (1990), "Modelling Price Behaviour in the European Car Market: 1970-1985," Economics of Industry Discussion Paper no. EI/1, STICERD, London School of Economics.

Nocke, V., (1998), "Underinvestment and Market Structure," Economics of Industry Discussion Paper no. EI/22, STICERD, London School of Economics.

Nocke, V., (1999), *Industry Structure and the Dynamics of Competition*, PhD Thesis, London School of Economics.

Oldale, A., (1997), *Dynamic Non-Price Strategy and Competition: Models of R&D, Advertising and Location*, PhD Thesis, London School of Economics.

Pakes, A., Gowrisankaran, G. and McGuire, P., (1993), "Code for Implementing the Pakes-McGuire Algorithm for Computing Markov Perfect Equilibrium." Mimeo, Yale University.

Raith, M., (1996), "Product Differentiation, Uncertainty and the Stability of Collusion," Economics of Industry Discussion Paper no. EI/16, STICERD, London School of Economics.

Robinson, W. and Chiang, J., (1996), "Are Sutton's Predictions Robust? Empirical Insights into Advertising, R&D and Concentration," *Journal of Industrial Economics*, vol.44: 389-408.

Scott-Morton, F, (1997), "Entry and Predation: British Shipping Cartels 1879-1929," *Journal of Economics and Management*, vol. 6(4): 679-724.

Shaked, A. and Sutton, J., (1983), "Natural Oligopolies," *Econometrica*, vol. 51, pp. 1469-1484.

Shaked, A. and Sutton, J., (1984a), "Natural Oligopolies and International Trade," in H. Kierzkowski (ed.) *Monopolistic Competition and International Trade*, Oxford: Oxford University Press.

Shaked, A. and Sutton, J., (1984b), "Product Differentiation and Industrial Structure," Theoretical Economics Discussion Paper no. 85/113, STICERD, London School of Economics.

Shaked, A. and Sutton, J., (1987), "Product Differentiation and Industrial Structure," *Journal of Industrial Economics*, vol. 36: 131-146.

Sutton, J., (1991), *Sunk Costs and Market Structure*, Cambridge, MA: MIT Press.

Sutton, J., (1998), *Technology and Market Structure: Theory and History*, Cambridge, MA: MIT Press.

Symeonides, G., (1997a), "Price Competition and Market Structure: The Impact of Restrictive Practices Legislation and Concentration in the UK," Economics of Industry Discussion Paper no. EI/18, STICERD, London School of Economics.

Symeonides, G., (1997b), "Cartel Policy, Non-Price Competition and Market Structure: Theory and Evidence from the UK," Economics of Industry Discussion Paper no. EI/19, STICERD, London School of Economics.

Tirole, J., (1990), *Theory of Industrial Organisation*, Cambridge, MA: MIT Press.

Walsh, P. and Whelan, C., (1998), "Modelling Firm Size Distributions in Food and Drink Products," Economics of Industry Discussion Paper no. EI/23, STICERD, London School of Economics.