A NOTE ON CONSUMER-PRODUCER COLLUSION IN THE PRIVATE MARKETING OF COLLECTIVE-TYPE GOODS

by

Earl A. Thompson
University of California, Los Angeles

Discussion Paper Number 142
January 1979

In a recent issue of this journal, Thomas Borcherding claims to produce a "counterexample" which "refutes" my theorem on the competitive oversupply of privately produced collective goods (1968). While his argument explicitly violates a condition of the theorem -- namely, the Cournot assumption -- and therefore is certainly not a counterexample, the alternative pricing contracts which form the basis of Borcherding's critique are well worth examining. For a further analysis of his "complex contracting," in revealing such contracting to imply all of the known properties of collusive contracting, helps extend our understanding of the nature and limitations of collusion. At the same time, it leads naturally to a general specification of the transaction costs implicit in any competitive model.

After discussing Borcherding's contracts and the general nature of collusion, I shall indicate how my competitive oversupply result holds under a substantial generalization of the Cournot-type information conditions in my previous papers. While this generalization will serve to sharpen the contrast between a collusive world and my own, more importantly, it will bring my competitive theory much closer to empirical applicability to observed markets for collective goods.

I. Borcherding's Contracts and the General Nature of Collusion.

Borcherding's argument is essentially that rational collective good sellers will individually commit themselves to a fixed price to any given buyer independent of the behavior of that buyer. Such price commitments clearly benefit the buyer in that he can then reduce his purchases from other sellers without suffering price increases from the committed sellers, such increases being the source of my collective goods overvaulation result. The contracting sellers, it is argued, are not hurt by such contracts because they receive essentially the same price with or without the price commitment.

However, since the only advantage of these fixed-price agreements is that the buyers making them can reduce their purchases from other producers without fear of price increases from the producers offering the fixed-price contract, it is easy to see how such agreements cannot be in the interest of a producer once we apply the Cournot assumption that the quantities sold of the other producers are fixed. For under a Cournot assumption, competitive bidding by the threatened producers would simply lower all prices and result in no ultimate benefit to the initiating producers.

Borcherding's attempted counterexample, relying as it does on the ability of the customers and some of their suppliers to gain at the expense of other suppliers while substantially altering the output choices of these other suppliers, strongly suggests that his "complex contracting" assumption requires collusive, cartel-type agreements. To isolate this issue, let us accept Borcherding's price contracts a priori, and let us further accept, at least temporarily, his conjecture that such pricing implies that a Lindahlian Equilibrium will be established under constant costs so that each customer pays a collective good price equal to his marginal real use value for the good.

Does the resulting equilibrium pass any of the familiar Stigler tests for the absence of collusion? First: If collusion is absent, there should be no incentive for a firm to secretly expand his output, joining together with some customers to upset the allocation. This test is not passed. It would pay an undetected producer to expand his output and sell an extra unit of the collective good for the market prices plus separate fees, and thereby lower the marginal use values -- and hence prices -- charged by the other producers.[1]

[1]Borcherding might be tempted to cut the knot here, replying that equilibrium use values are known a priori by the contracting parties and the contracted prices are simply fixed at those levels. But if this were so, then everyone would know entire the solution a priori and there would be no point to a market

The existing firms would, if they could not detect these side contracts, recognize the weakness of their initial positions and charge higher initial prices in order to prevent the losses that would otherwise result from such contracts. It would then be but a short step to my overproduction equilibrium. But since we are maintaining Borcherding's conjectured equilibrium, the side deals must be detectable and each consumer must be pre-committed to the existing sellers to refuse all such temptations.

This means that Borcherding's contracts are tantamount to multilateral contracts: They must be agreements contingent on transaction with third parties. Such agreements, when they involve all possible producers, as they must in this case, have the same, extreme, informational conditions of effective cartel agreements between otherwise competitive producers. For such cartels require at least one of the colluding parties to know the transactions of all of the others with outside parties and to effectively communicate sanctions on numerous such transactions. In other words, once we recognize that the costs of making contracts containing contingencies on the behavior of outside parties are prohibitive in competitive general equilibrium theory (Thompson, 1972), then Borcherding's theory immediately fails as a competitive general equilibrium theory.[2]

1 (con't)
mechanism in the first place. The point is that we must assume that the buyers are the only ones who know their own, entire preference functions even though it is necessary for any kind of perfectly discriminating price system that the sellers know the marginal benefits of the buyers once a solution is reached. The latter information assumption is acceptable for decentralized, market analysis while the former is not. Nevertheless, Borcherding's equilibrium, as we are about to see, can also appear under the more realistic informational conditions acceptable for decentralized market analysis. Our analysis throughout the text of this paper maintains these latter relatively realistic, informational conditions.

[2]This test for being as competitive general equilibrium model is unnecessarily stringent in the sense that we can imagine competitive worlds with some contracts conditional on the behavior of third parties even though existing formal models of competitive general equilibrium obviously contain no market or asset-accounting for these contracts. Nevertheless, when we allow zero costs of all such conditional transactions, the model necessarily degenerates into one containing a purely monopolistic market structure. (Thompson, 1972).

A second Stiglerian test for the presence of a competitive equilibrium is the absence of any incentive toward non-contractual price changes. Borcherding's hypothesized equilibrium also fails this test, for each producer in his world loses no sales by raising somewhat the non-contracted price of his product. While Borcherding's buyers have contracts assuring them of given nominal prices from the sellers, the latter can threaten to cut a buyer off if he refuses to pay the non-contractual charges (e.g., if he refuses to pay the transportation or transmission costs). For if the buyer were cut-off, he would have a still higher marginal real value and therefore have to face higher effective prices from the other sellers.

Summarizing, while perfectly competitive equilibria imply the absence of incentives toward production increases with side deals and non-contractual price changes, the substitute, Lindahlian equilibrium described by Borcherding implies both of these incentives. Borcherding's substitute equilibrium requires contracting costs which would imply cartels in ordinary private-goods theory. Alternatively, Borcherding's "complex contracting" requires sellers to know, and make contracts contingent on, the transactions of each of the many buyers with other sellers and therefore implies a level of contract information which is too perfect to support the competitive marketing of <u>any</u> good.

All this relates back to the Nash-Cournot assumption. When parties contract conditionally on the behavior of third parties, it is to alter the strategies of the latter parties. Since a Nash-Cournot assumption assumes that the strategies of third parties are given, allowing conditional contracts means dropping the Cournot assumption. The problem with

the constructive part of Borcherding's discussion is that he did not replace my Cournot assumption with an alternative interaction assumption and therefore did not produce even an outline of a complete model. Since his analysis could not yield a determinate solution he was reduced to conjecture as to the nature and ultimate efficiency of his "complex contracting" interaction.[3]

Borcherding's failure to see that his constructive price function analysis was incomplete is apparently based on his failure to appreciate that the basic theoretical problem requires deriving what economists usually accept as obvious -- the price functions which constrain individual behavior in a given market structure -- from a prior theory of economic interaction requiring some game-theoretic specification (such as a Cournot assumption) to represent each party's information structure. Rather than producing such a theory, Borcherding is content with a traditional methodology which eschews game theoretic assumptions and argues without any apparent theoretical discipline what "competitive price functions should look like." While I was so indirect in indicating this theoretical requirement that Borcherding cannot really be blamed for employing the traditional methodology, all of the alternative derivations of my overproduction result which I am aware of [viz., Auster (1978), Compton (1976), and Heiner (1975)] have collective-good price functions which are developed from well-specified informational conditions.[4]

---

[3] My '72 paper outlined a solution (which was extremely robber-barronish) under an interactive assumption appropriate to perfect contract information (including conditional contracts), which has since been generalized (Thompson-Faith, 1978b). Under these extreme, perfect information assumptions, an equilibrium is almost always a Pareto-optimum, as Borcherding conjectures.

[4] The latter two substantially generalize my model in such a way that competitive price functions explicitly result from restrictions on the information structure similar to those outlined above. Borcherding was apparently not aware of any of these interesting theoretical studies.

Nonetheless, it is clear that Borcherding's contracts imply a buyer-seller cooperative cartel. If the cooperating group is not the entire population, it will take advantage of its effect on prices to achieve a redistribution from outsiders, obvious monopoly-type inefficiencies will result, and Borcherding's optimality conjecture is false. If the cooperating group is the entire population, then the cooperating group can potentially achieve an optimum in that it internalizes all effects. However, Borcherding's price rule -- charging everyone a constant collective good price equal to his use value of the collective good -- is essentially a simple Lindahl rule and has the same defect unless there is perfect knowledge either of changes in each buyer's behavior or of the entire social optimum (in addition to the relatively reasonable perfection of knowledge of the buyers' equilibrium use values of the collective good). That is, as we have been pointing out, absent such relatively extreme informational perfection, any buyer can benefit by altering his behavior so as to alter his marginal use value in a solution and therefore alter his price for this good.

Consequently, when the entire group is in the cartel, it is not surprising that Borcherding-Lindahl pricing is likely to be dominated by other, more coercive, political allocation devices. While this is not the place to discuss such devices, I must say that I agree with Borcherding's general position that cooperation largely solves the competitive overproduction of public goods. In fact, I used this hypothesis in my '68 paper to rationalize some otherwise anomolous features of observed governmental policing with respect to collective goods. Our only disagreement in this area is how such cooperative works. I think I see it working through coercive, political processes and he conjectures

that it works through voluntaristic, private behavior. The only defects in his argument for this position are (1) his allegation that this private, voluntaristic cooperation is consistent with competitive markets and (2) his apparent lack of recognition of the additional, extreme, informational conditions implied by such cooperation.

The question that remains is how Borcherding -- who did recognize that he was not using a Cournot assumption and that his buyer-seller contracts implied incentives toward the monopsonistic under-purchase of inputs into the production of collective goods in the increasing cost case (implicitly when the group did not contain the entire population) -- was able to convince himself that his contracting was of a competitive variety. My best answer is that he apparently felt that my world was even less competitive than his. For his critique emphasizes my subsequent restatement and generalization of the model in which I analyze the microeconomics of the buyer-seller relationship for collective goods as a bilateral monopoly relationship. And he concludes that my model "*describes the behavior of an open but perfectly price discriminating cartel*" [his italics]. But he fails to note that my bilateral monopoly is a natural one that necessarily applies to any world -- including his own -- in which collective goods are privately sold and that the familiar competition vs. monopoly distinction arises at a less microeconomic level, where the issue is whether or not there are many contractually independent sellers of the substitute products. Since my competitive world has the sellers contractually independent while his world does not, his is clearly of the monopolistic variety while mine is of the competitive one. And if my world describes

the behavior of an "open cartel," Borcherding's is describing the behavior of a "closed cartel" and therefore remains less competitive even by this criterion.[5]

II. Generalizations

Having referred to my representation of any collective-good buyer-seller relationship as one of natural bilateral monopoly, Borcherding could have easily extended his critique to the corresponding generalization of my over-valuation result in which the victims of any perfectly discriminating monopolistic will over-buy from producers of substitutes. This generalized inefficiency, like the special case of collective goods overvaluation, does not hold when the discriminator can observe each of his customers purchases of all other goods and make contracts contingent upon all such purchases. But, as above, such a "refutation" of these generalized inefficiency results is tantamount to allowing perfect, n-person, cooperative interaction and is way beyond the intentions of decentralized market analysis.[6]

---

[5]Nevertheless, I strongly resist the "open cartel" characterization of my model, even though Borcherding's claim is only that the "behavior" is the same in my model as one with an open cartel. The reason is simply that my model implies no cartel-like agreements and a private-good, open cartel is known to produce essentially competitive results anyways.

[6]Nevertheless, the received economic theory of a single perfectly discriminating monopolist implicitly endows the seller with these extreme, cooperative-type, informational powers. Even here, however, the implicit, extreme, informational perfection, and the corresponding efficiency result, necessarily break down when there are two independent, perfectly discriminatory monopolists. (Thompson, '69, '73). This is because treating the two monopolists as "separate," or "independent," producers implies a non-cooperative interaction between them and hence prohibitively costly contracting contingent on the purchases from one another. The result is that the general over- and under-valuations described in the text above occur, but only with respect to the purchases of outputs of the discriminating monopolists. When such perfect discriminators are numerous, their respective outputs are related by perfect substitutability, and the outputs are inelastically supplied to each customer, it is easy to see (Thompson, '69, '73)

-continued-

My essential conclusion should, however, hold under some, substantive generalization of my Cournot-Nash assumption. While the standard, all-private-goods, competitive model can be approximated by a many-firm model of rational interaction under the Cournot-Nash assumption of zero information regarding third party actions, it is also true that the standard model is approximated by a many-firm model of rational interaction under a Stackelberg-Von Neumann-Morgenstern assumption assumption of "perfect information" regarding third party actions. (For a proof, see Thompson-Faith, 1978a). It would appear that any informational perfection short of the communication of strategies contingent on the behavior of third parties will produce a familiar type of value theory for private goods. (Perfect strategic communication, which implies contracts contingent on third-party behavior, produces a theory unlike anything in standard value theory and does not approximate a standard competitive model even when the number of firms is large (Thompson, 1972, Thompson-Faith, 1978a and 1978b).)

So my general value theory for perfect discriminators, and its application to the case of competitively marketed collective goods, should also maintain its major features when we replace the Cournot-Nash assumption with one which

---

6 (con't)
that we have a special case of a competitively marketed, pure collective good. The only difference between the resulting model and my '68 model is that the former allows the collective good sellers to make price contracts contingent on customer purchases of non-collective goods and hence the overvaluation of private-good substitutes (and undervaluation of private-good complements) disappears while the overvaluation of collective goods remains. (This is formally established in the elegant general equilibrium model of Heiner (1975)). Nevertheless, the perfectly cooperative interaction implicit in part of this model certainly prevents it from describing a uniformly competitive world.

allows discriminators to observe -- but not make contracts contingent on -- their customer's purchases of other goods. Again, the natural alternative, short of strategic communication, is the polar, "perfect information" assumption of Von Neumann-Morgenstern and Stackelberg in which any discriminating seller can costlessly observe each of his customer's prior purchases and perfectly predict his subsequent, rational purchases. What does it imply?

Consider first the case in which the purchase and use of a given substitute or complement precedes any communication between a perfect discriminator and his customers. Then, even if the discriminator is able to observe these actions of his customers, he must take them as fait accompli when he finally deals with his customers. So the customers will clearly have the same incentive to overbuy substitutes (and underbuy complements) in anticipation of dealing with the perfect discriminator that he has under a Cournot-Nash assumption.[7] While this incentive would disappear for substitutes and complements purchased subsequent to the purchase from the discriminator, a temporal recurrence of purchases from future discriminators would obviously reintroduce the same general effect so that roughly the same inefficiency would result.

In the special case of a series of sellers of different units of single collective good, the analysis is similar: Assume first that the collective good is purchased once and for all. The first seller is able to sell his output at prices far in excess of the customers' real use values of the good

[7] The work of Batchelder on compensation rules under eminent domain provides an example of this inefficiency. A property owner, who anticipates that he may lose his property to a government which will only pay what the property is worth in the private sector, will overvalue current investment activities which increase the property's value to the private sector and undervalue those which increase the property's value to the public sector.

because having an extra unit of the good will lower their use values, hence lower their purchase prices, from all of the future sellers. The same is true for the second seller except that the first seller has already made his deal so that this one set of prices is not reduced. This continues on until the last seller in line just receives the customers' real use values for the good. Since the outputs of all firms (except one) exceeds that at which the sum of the marginal use values exceed the marginal cost of producing the output, there is overproduction by all firms (except one). If we now allow recurrent chains of collective-good purchases, then even the last seller within the period will be able to charge significantly in excess of the customer's real use value for the collective good.

The asymmetry and relative complexity of this model, and the fact that its welfare results are qualitatively identical to those under a Nash-Cournot assumption, led me to choose the latter assumption. Nevertheless, the most important privately marketed collective good observed in the real world, industrial technology, is marketed in a sequential fashion reflecting a Stackelberg-Von Neumann-Morgenstern rather than a Cournot-Nash interaction. That is, the typical innovation is introduced at substantially different times by different firms, each of which is generally informed about the experiences of prior innovators and can infer a fairly predictable pattern of future responses from past behavior (see, e.g., Mansfield).

At least since the writings of Schumpeter, economists describing this innovation-diffusion process have inferred from the relative infrequency of observed patent licensing agreements that the firms which follow others in adopting a new technology are typically able to free-ride on the knowledge gained by the prior innovators. However, an alternative inference is that such firms must pay for a new technology by hiring quasi-monopolistic technologists who

have previously acquired information required for the practical application of the new technology. In this view, an innovator may, in effect, collect close to the value of his information from every firm, in which case the conditions of our generalized model for the competitive marketing of collective goods would apply. The first firm to which an innovator sells (usually the firm which has financed the basic research and has the highest demand price for the research output) would pay too much for his services to the extent that having them would lower the future prices paid to other technologists. Since the first buyer of the technology would enjoy a permanent monopoly on its use if the information did not get out, it is in the joint interest of the innovator and the firm to keep the technology to themselves. But this is seldom observed. For it is also in the private interest of the innovator, who usually receives only a small fraction of the joint monopoly return, to arrange side-deals with the firm's competitors or other technologists who will be able to sell to these competitors.[8] So these particular buyer-seller collusions usually break down. Recognizing this, the buyer is willing to pay little more than his competitive use value of the innovation, in addition, of course, to its value in changing the prices he must pay to sellers of other technologies.

So, after selling his services to the firm with the highest demand price for his innovation, the innovator will then find it profitable to sell his idea to the next highest demander of the new technology in another bilateral monopolistic interaction. To do so, the innovator either quits the first

---

[8] In case the innovator could bargain close to the entire return from the innovation to the firm, these side-deals would not substantially pay. However, it would then generally pay the firm to hire additional innovators to prevent such "hold ups" so that the eventual dissemination of the information to other firms would still by very likely. For since the firm would pay no more than the entire (monopoly) value of the innovation to the innovation team, any individual member, who receives only a small fraction of the value, would generally profit by selling the innovation to a similar firm for close to the entire (duopoly) value.

firm and goes to work for the second or trades his knowledge to another technologist (usually a co-worker) who then goes to work for the other firm.[9]

Since a competitive firm does not gain substantially from the existence of a new technology (as contrasted with the firm's use of the technology given its existence) many writers have argued that expenditures on basic research by rational firms implies at least short term expected monopoly returns[10] from

---

[9]To the extent that the first firm can acquire a patent right to the innovation, the innovator will have to take a lower price for his services because of the relatively high cost of the second firm's innovating around the patent. This will increase the temporary monopoly gain to the first firm as it will slow up the diffusion of the idea to his competition. While this suggests that patents are merely devices to increase the length of the temporary monopoly resulting from an innovation, in a more realistic environment where some ideas can be figured out by outside innovators without the help of the original innovators, the patent right serves to provide sufficient return to innovation to arrive at a second-best social optimum despite the extension of the temporary monopoly firm supporting the original research.

[10]An additional possible monopoly return from innovating, pointed out by Hirshleifer, is that an innovator may profit from an hypothesized monopoly over the knowledge that his new technology exists by making speculative transactions in private-goods markets prior to the acquisition of such knowledge by others. The problem with this argument is that when an innovator has a monopoly over the knowledge that the new technology exists, it pays others to monitor his speculative orders. Other speculators will not transact with him, and, even if he could execute some orders by finding some transactors who do not know who they are trading with, numerous other speculators would duplicate his trading behavior and reduce his profits to near zero. While such speculation perfections are not empirically observed, neither do we observe innovators engaging in substantial trading on their information. This is likely because they lack the other information necessary to be successful speculators. There is also an apparent absence of substantial sales of information regarding the existence of innovations by innovators to professional speculators. This is likely due to the fact that any attempt at acquiring a substantial, non-monopsonistic, price for the information about the new technology from market experts would reduce the value of the information to insignificance.

the research. But this is an incorrect inference. The innovator who could not find a buying firm to finance his research would likely go elsewhere to finance it and charge the buying firm which would otherwise finance the research a much higher price for the eventual innovation. Therefore, the financing of research by the buying firm may have no effect on the quantity of research done; the firm may support the research (in exchange for rights to use any eventual research product) simply in order to reduce the price the firm would otherwise pay for the innovator's output. The obvious gain to this vertically integrated arrangement in terms of saving the transaction costs of selling an idea for the first time would be a plausible rationalization for the observed relative scarcity of the specialized research firms.

The fact that the firm which supports the basic research and is therefore first to have the technology available to it typically pays much more than later firms in order to use the research, is evidence in favor of our hypothesis that the first firm to purchase an innovation has a substantially higher value and price than subsequent purchasers. Further evidence is that by far the most important determinant of the order in which firms adopt an innovation is the order of the total profitability of the innovation to the firms (Mansfield, et al.).

A direct empirical implication of the entire theory is that there will be an abnormally high labor mobility among technologists. An implication of the related overvaluation of innovations is that firms will acquire an abnormal "excess capacity" of rights to use various technologies. Professor Steve Cheung, who has been studying the institutional characteristics of the research and development industry for several years now, informs me that the two most unusual features of the industry are (1) the high mobility of the

technologists and (2) the high stock of unemployed technologies acquired by the buyers.[11]

Since the above tendency to overvalue innovations occurs because of the effect which an innovation will have on the firm's purchase prices for future innovations, and since it is obviously highly impractical for future innovators to precommit themselves to fixed prices for their future innovations so as to prevent such current overvaluations, Borcherding's collusive pricing agreements would be highly impractical for this application. Nevertheless, another kind of collusive interaction, one between current innovators, often occurs and does not meet with the kind of government interference which prevents collusions in other sectors of our economy (Thompson-Faith, 1978a). Innovating firms will often "stake out" an area of research or development and punish would-be competitors by a predatory "raiding" of their companies for top technical personnel to prevent wasteful competition in the same line of research or development. While such commitments to respond to the actions of others are of the same logical form as collusive contracts, they deal only with current and near-future interactions among innovators and therefore have much less severe informational conditions than Borcherding's contracts. Moreover, the welfare implications of such commitments are much different. For without such commitments, there would be an inefficient "rush to invent" (Barzel) because there are no governmentally preassigned rights to produce

---

[11] Professor Cheung's main reference on the unemployed technology data is a study by Gharrity. He did, however, express the opinion that these studies overstate the stock of unused technologies. His main reference on the high mobility of technological personnel was an observation that 95% of an apparently random sample of research directors of large companies during a recent, 3-1/2 year, period had switched employers.

certain, prespecified, types of innovations in the real world. One obtains a right to an idea by producing it before anyone else does. By allowing innovators to stake out certain research areas with predatory, or collusive, reaction functions, the government greatly reduces this tendency towards excessive impatience in the development of an idea. So such behavior works merely toward an efficient definition of property rights and permits us to use a simple property rights model, such as my competitive model, to describe the resulting allocations.

As emphasized in my '68 paper, once we introduce the realistic informational imperfections in which technology sellers cannot collect the entire values of the technology to the various buyers and cannot practically exclude several non-paying users of the technology, the competitive outputs of the collective goods come down to levels which cannot be said with confidence to represent either under- or over-production of innovations. (A recent paper of mine indicating substantial statistical support for the approximate Pareto optimality of the level of U.S. innovation will be sent on request.) Furthermore, once we combine the imperfection in the ability to exclude non-paying users of a technology with realistic imperfections in the abilities of firms to "stake out" certain areas of research, the "rush to invent" tendency resulting from the latter imperfection is countered by a "wait-for-the-other-guy-to-invent-it-first" tendency fostered by the former so that the direction of the probable mistiming of innovations cannot be identified a priori either. (For a similar, more formal development of this point, see Kamien and Schwartz.)

In conclusion, while the most empirically important market for collective goods, the observed market for technical innovations, is probably much closer

to a competitive variety than it is to Borcherding's collusive form, much more careful empirical work is required before we can make confident statements regarding the empirical power of a simple competitive model or the under- or over-production of innovations in a realistic setting.

REFERENCES

Auster, R.D., "Private Markets in Public Goods," QJE, August 1977, pp. 419-430.

Barzel, Y., "The Optimal Timing of Innovations, Rev. Econ. & Stat., August 1968, pp. 348-355.

Batchelder, Ronald W., "Optimal Compensation Under Government-Imposed Rearrangements of Owernship," UCLA Ph.D. dissertation, 1975.

Borcherding, T.E., "Competition, Exclusion, and the Optimal Supply of Public Goods." J. Law & Econ., April 1978.

Compton, G.R., "The Economics of Discretionary Collective Goods," UCLA Ph.D. dissertation, 1976.

Gharrity, N.J., "The Use and Non-use of Patented Innovations," Ph.D. dissertation, Johns Hopkins University, 1965, esp. p. 265.

Heiner, R.A., "The Exchange of Collective Goods and the Existence of Generalized Competitive Equilibrium," UCLA Ph.D. dissertation, 1975.

Hirshleifer, J., "The Private and Social Value of Information and the Reward to Inventive Activity," Amer. Econ. Rev., Sept. 1971, pp. 561-74.

Kamien, M.I., and N.L. Schwartz, "The Timing of Innovations Under Rivalry," Econometrica, Jan. 1972, pp. 143-60.

Mansfield, E., The Economics of Technological Change, New York, Norton, 1968.

_____, et al., The Production and Application of New Industrial Technology, New York, Norton, 1977, Ch. 6.

Schumpeter, J. A., Business Cycles, New York, McGraw-Hill, 1939.

Stigler, G. J., "A Theory of Oligopoly," J. Polit. Econ,, February, 1964.

Thompson, E. A., "The Perfectly Competitive Production of Collective Goods," Rev. Econ. & Stat., February, 1968, pp. 1-12.

_____, "The Perfectly Competitive Production of Collective Goods: Reply," Rev. Econ & Stat., November, 1969, pp. 479-482.

_____, "A Reformulation of Orthodox Value Theory Containing an Explicit Information Structure," 1972, UCLA Working Paper #19, 73 pages.

_____, "The Private Production of Public Goods: A Comment, J. Law and Econ., October 1973, pp. 407-412.

_____ and Faith, R. L., "A Model of Rational Noncompetitive Interdependence," UCLA Working Paper #86, 1978a.

_____ and _____, "A Theory of Games with Truly Perfect Information," UCLA Working Paper #61, 1978b.