PEPPERDINE UNIVERSITY

The Journal of Entrepreneurial Finance

Volume 2
Issue 2 *Spring 1993*

Article 2

12-1993

Innovation in Large and Small Firms

Kenneth J. Arrow
Stanford University

Recommended Citation
Arrow, Kenneth J. (1993) "Innovation in Large and Small Firms," *Journal of Small Business Finance*: Vol. 2: Iss. 2, pp. 111-124.
DOI: https://doi.org/10.57229/2373-1761.1159
Available at: https://digitalcommons.pepperdine.edu/jef/vol2/iss2/2

Innovation in Large and Small Firms

Kenneth J. Arrow

I. INTRODUCTION

This essay is intended to begin the elaboration of a theme: the interaction between the observed sizes of firms and their internal decision making procedures. This theme is a major one in the symphony of entrepreneurial activity. The entrepreneur, as the maker and changer of economic and productive life, is usually envisaged as an individual. In the neoclassical tradition, he (or, rarely, she) is the lightning calculator, the individual who rapidly scans the field of alternative productive processes and chooses the optimum at any given set of prices. In the Austrian tradition, most notably in the work of Schumpeter,[1] he is endowed with a special psychology that makes him all the more an individual in the strict sense of the word; he cannot be replaced by a machine or by a multiplicity of individuals, who would inevitably slow him down. "He travels fastest who travels alone," says an ancient proverb.

However, the individual entrepreneur-proprietor does not loom nearly as large today as suggested by these accounts. The large, even giant, firm is a massive presence on the economic landscape. These large firms not only predominate in the static allocation of resources, but are the sources of much of the world's change. They share fully with others as the sources and users of innovations.

This is not to deny the continued importance of the relatively small firm and the individual inventor. Indeed, the coexistence of large and small firms is itself an interesting intellectual question. If, in fact, large firms do have advantages over smaller ones, why are small firms not eliminated in the competitive struggle? More generally, if there are differential advantages to one size or another of firm, why do firms not converge to the optimal size?

The presence of large firms creates logical difficulties for the concept of property and for the reward structure of the individual, as Berle and Means[2] pointed out almost fifty years ago. The sharp calculating eye of the

This article has been reprinted with permission from Professor Kenneth J. Arrow. It was first published in *Entrepreneurship*, Joshua Ronen (ed.), 1983; Published by Lexington Books, D.C. Heath and Company, Lexington, MA.

Kenneth J. Arrow • Department of Economics, Stanford University, Stanford, CA 94305.

The Journal of Small Business Finance, 2(2): 111-124
ISSN: 1057-2287

neoclassical entrepreneur was for his own profits, and even those who gave a more psychological interpretation to entrepreneurial motives could hardly deny that revenue was essential among them. But an employee, however entrepreneurial in spirit, does not have property rights and cannot claim profits, the residual revenues after contractual claims. Much ingenuity can go into alternative compensation schemes, but the maker of decisions about innovations can no longer be simply identified with the recipient of rewards (and taker of losses) from them.

Of equal, or even greater, significance is the diffuse control structure of the large firm. Essentially, no one can make decisions without limits even within the framework of feasibility. Even a chief executive officer is restricted, partly because of the need to adhere to well-defined operating procedures, and partly because limits on span of control prevent him from making more than a limited range of decisions with limited information. In large firms, entrepreneurship has sociological as well as psychological and economic dimensions.

The remarks thus far show that entrepreneurial activity, however defined, operates in different ways in large firms than in small ones. I will concentrate here on entrepreneurship as Schumpeter conceived it—the process of innovation. The basic decisions are the recognition of promising ideas and the financing of their development. We want to discuss how these decisions operate in firms of varying sizes.

An economist would not, of course, discuss any issues of decision making by firms without taking account of market relations. Since the development of innovations is an investment, the most relevant market is the capital market. However, innovations are, by their nature, rather odd commodities from a neoclassical viewpoint; they tend to be indivisible. Their development is attended by uncertainty; if everything about an innovation were known, it would not be an innovation. What is still more, the properties and economic potential of an innovation are by its nature likely to be better known to the innovator than to a prospective source of financing.

In short, the supply of capital for innovation is not modeled well by conventional competitive market theory. Indeed, most of the analysis in this chapter will center about the methods of financing innovation and their implications.

The chapter is organized as follows: In the next section, I review in the sketchiest way the idea that large firms are really significant in the economy and constitute a phenomenon about which we cannot be indifferent. I then describe an idealized model of the process of innovation (oversimplified of course) designed to serve as a basis for subsequent discussion. The heart of the paper follows, an attempt to understand the factors in the decision to innovate and (what is essentially equivalent) the financing of the innovative

activity. In particular, I stress the systematic variation of these decisions with firm size and complexity. It is concluded that there is likely to be tendency toward specialization, less costly and more original innovations will come from small firms, and those involving higher development costs but less radical departures in principle will come from larger firms. This specialization creates opportunities for trade, as all specialization does; in this case, the trade will frequently be in firms as such; that is, takeovers and mergers.

II. THE SIGNIFICANCE OF LARGE FIRMS

From the popular viewpoint, the concentration of economic power is one of the most obvious aspects of the economic world. In mainstream neoclassical economics, it hardly appears, especially in more abstract versions of the neoclassical system (for instance, Arrow and Hahn).[3] Of course, the presence of natural monopoly is recognized, and this is the basis for the doctrine of price regulation. Even here, many economists consider that there is sufficient competition from substitute products to make natural monopoly an unimportant concept.

The trouble is that the analytic tools of neoclassical economics are not well adapted to departures from perfect competition. There are two pillars to the edifice: the optimizing behavior of the firm and household, and the equilibrating forces of the markets that link them. Optimizing behavior can indeed be discussed under conditions of market power; the theory of monopoly is rich in implications. But the concept of imperfectly competitive markets is very hard to define. Various ad hoc constructions, such as Chamberlin's notion of monopolistic competition[4], have appeared, but they suffer from inconsistencies. Game theory has supplied a formal framework that, in principle, replaces markets by more general forms of interactions, but it has not yet succeeded in producing a general theory comparable in power to the theory of general competitive equilibrium.

Hence, there is a bias toward analyzing the competitive case. As we know, this analysis requires, if taken literally, that the production possibility sets be convex. In particular, it requires constant or diminishing returns to scale. The latter case suggests a bias toward small firms; under free entry, the smaller the better. Constant returns, on the other hand, are neutral toward the size of firms. If two firms merge, the owners will (under perfect competition) be neither better off nor worse off than they were before. Under perfect markets, including perfect capital markets, the profits of two different activities will simply be additive.

A good deal of the empirical literature on firm sizes has been devoted to arguing that the competitive model is adequate in practice; that is, there are not many markets dominated by one or two firms. This may well be true; it follows that the static efficiency characteristics of competitive equilibrium can be postulated to hold in the real world. This is very far from denying the existence of very large firms or from explaining this phenomenon.

For it is certainly a fact. Depending on what measures you use, 500 firms constitute half or more of the nonagricultural economy.[5] It is frequently argued that the indices of concentration have not shown much secular rise, at least not for 75 years or so. However, this misses the point. The economy has grown enormously in this period. If it were merely a question of replication, that is, if the economy were expanding homogeneously, we would expect the number of firms to increase in the same proportion. Since firms differ in size, for whatever reason, we would expect the proportion of firms of a given size to be constant, while the total number increases.

To be sure, the expansion of the economy has not been merely a replication. The fact that per capita income is rising and, more strongly, that factor productivity (output per unit input) is or was rising, implies a change in the proportions of the economy. But one component of growth in the market remains sheer population-size or total factor supply (capital and labor). One might expect then, that the number of firms would be proportional to the extensive growth of the economy (its size in population or inputs), while the size of each firm might be expected to grow with the intensive growth (for example, output per capital or per unit factor supply). This is what has happened (roughly) to the distribution of individual income. It can be expressed as the constancy of the Lorenz curve; that is, the proportion of total income received by a given proportion of the population arrayed by income level (for example, the upper tenth of income recipients) is a constant.

But this is not what has happened with the distribution of firm sizes. The proportion of total sales or income received by a fixed number of firms has more or less remained constant.[6] Therefore, the proportion of income received by a fixed proportion of firms starting from the top (say, an upper decile again) has increased.

In short, we find that the size of each firm has increased more than proportionally to intensive growth. If intensive growth is identified with productivity (either of labor or of total inputs), it follows that not only the outputs but also the inputs of the average firm have risen. This implies that the forces determining the sizes of firms (in particular the economies of scale and the size of the market) have so shifted as to make larger firms more advantageous.

The increasing costs of innovation are a possible candidate, and the later analysis in this chapter implicitly makes the case for this proposition. At

this point, however, I only wish to establish that there has been a significant shift to larger sizes of firms and that this shift has systematic economic consequences and causes.

One obvious feature of larger firms (as contrasted with smaller contemporary firms or even with the same firms when they were smaller) is that they are more complex. They are not simply scale expansions of smaller firms, any more than the economy as a whole is a scale expansion of its earlier historical self. Even if the added activities are similar in nature to original ones, random fluctuations would make coordinating activities profitable. More broadly, growth usually involves disaggregation of activities and differentiation of products and activities.[7] No doubt these tendencies can ultimately be explained in terms of indivisibilities and other causes of increasing returns to scale. The complexity requires additional control functions at the central level.[8]

Coordinating activities themselves are costly; not only do they directly involve the use of resources (managerial and supporting personnel, associated equipment, space, and communication channels), but they also impose costs upon decision making at lower levels by creating delays and requiring additional communication costs. They are undertaken because the costs of coordination are exceeded by the benefits.[9] As Coase has argued, these benefits are relevant only if they are not obtainable by coordinating separate activities in the marketplace, through prices.[10]

This point can be emphasized by considering the multidivisional firm and the role of transfer pricing. A large firm is organized into profit centers, each of which operates as virtually a separate firm. Transactions between them are market transactions, and payments between them are made at current market prices or (if no suitable market exists) at transfer prices mimicking market prices. Presumably the opportunities for direct (as opposed to market oriented) coordination of activities have been exhausted within the profit centers. What distinguishes the large firm, however, from a collection of smaller firms is that many resource-allocation decisions are still make at a central level, particularly capital formation. A profit center is responsible for its own decisions on current flows, but in general it cannot make its own investment decisions, except possibly for very trivial ones. Indeed, it is surprising how often decisions on investment require the approval of the Board of Directors, while decisions of at least equal importance relating to pricing and production are decentralized to much lower levels.

There is, in short, an internal centralized mechanism for allocating available investment funds to specific projects among the various profit centers. The internal capital-allocation mechanism is not, properly speaking, a market; that is, a profit center cannot borrow any amount at a fixed rate

of interest. Rather, the project it proposes must be examined by the allocating authority for feasibility and profitability.

It would not be correct to contrast this allocation mechanism with an external capital market thought of as a true market in the textbook sense. Much, though not all, external financing is also project-specific and rationed. A bank does not lend by buying securities from anonymous sellers, but by lending to particular firms and individuals and often by looking at the particular project that the lender wishes to finance.

(We will not here study why capital allocation is so largely centralized, even in an otherwise decentralized firm. Part of the reason, certainly, must be the relatively slow feedback. The head of a profit center is not personally liable for the costs of bad investments. Considering job mobility, he may not be around to take any consequences when an investment is realized.)

It is important to distinguish the existence or absence of an internal capital allocation (in this sense) from the presence or absence of external financing. The supply of capital available for internal allocation can come either from retained profits in the various profit centers or from the outside. Large firms in general have an advantage in access to the outside capital market. One reason is the principle of insurance. Investing as they do in a variety of projects, their earnings are apt to be more stable and, therefore, the riskiness of their securities is reduced. Another reason is an economy of scale in attention and information-gathering from the viewpoint of the suppliers of funds. A large firm is a greater demander, and it therefore pays potential investors to concentrate their attention on that firm's activities rather than scattering it in one-shot transactions over many firms, for each of which there will be relatively little opportunity to use the information.

Diversification of activities also implies a more stable source of internal funds. Hence, in general, large firms will have a disproportionately larger and more stable internal capital supply than smaller firms will.

III. A MODEL OF THE INNOVATION PROCESS

Innovations are infinitely variable; indeed, they include all alterations in knowledge of current production relations between inputs and outputs. Most are very small, but those are not the ones we are concerned with here. We wish to stress those large enough that deliberate decisions are needed to proceed along the path of innovation. An innovation may never be realized as a product; if it is so realized, it may not remain in production very long. The process of innovation is, virtually by definition, filled with uncertainty; it is a journey of exploration into a strange land.

We take as a primitive of the system a stream of concepts—ideas for innovation. These occur to individuals both within firms and elsewhere. A concept may or may not prove to be feasible. If it is feasible, it may or may not prove to be profitable. These determinations require investment, and it is these investment decisions that we are investigating.

For simplicity, we will distinguish two further stages after the concept, those of research and of development. Somewhat arbitrarily, we will think of research as determining the feasibility of the concept and development as determining its profitability.

In this model, the concepts are random events, not controllable and unaffected by policy. They will, of course, depend on many factors, but especially on the state of knowledge in the relevant specialty. This in turn, may be influenced by previous innovations in the same intellectual area.

Decisions are made at the next two stages. First, research is needed to determine if the concept can be translated into actuality. The research may be more or less costly to carry out. When it is completed, it yields information about the prospects for development in the following sense. At the start of the development process, there will be a relation between the profitability of the innovation when it is finally introduced (possibly zero or even negative) and the amount invested in development. This relation depends on the information gathered in the research phase;that is, the profitability of the innovation at any given level of development expenditure with vary will the information obtained from research. Further, the relation of profitability to development cost is uncertain even given the research outcome.

To put it in a slightly different language, the research outcome is purchased by the research expenditure. The profitability of innovation after development is a random variable with a probability distribution conditional on both the research outcome and the decision about development expenditure. Given this distribution, the firm has the problem of choosing the optimal development expenditure. The optimal level might be (and frequently is) zero; but if it is not zero, it is frequently a very large amount.

In this simple model, there are two points at which decisions are made: (1) to engage in research and (2) to determine the optimal level of development expenditure. In the first decision, the information potentially available to the decision makers consists of the concept and publicly available information. In the second, it consists of the concept and the research outcome together with publicly available information. Of course, this oversimplifies the process in many ways. The sharp distinction between research and development is overstated. Furthermore, the development process itself is sequential. Instead of a single decision establishing development expenditures for an entire project, there are repeated

reassessments based on information revealed by the development process itself. However, our simplification will be adequate for our purposes here.

What must be insisted upon is the privacy of the information and its relation to the locus of decision making. The two relevant pieces of information, to repeat, are the initial concept and the research outcome. They are received in the first instance by some particular individuals. If these individuals were the decision makers, there would be no difficulty in principle. The decisions made (to engage in research and to choose the optimal level of development expenditures) would be optimal given the information available.

But the individuals concerned are members of organizations, small or large. The decisions to be made involve the allocation of resources. Some of these decisions might be structurally delegated to them. However, as the amounts involved increase, there will be more and more need for approval at higher levels. The internal capital-allocation mechanism will become involved. Those in the lower levels who have the relevant information cannot make the final decisions. Their scope of authority is often restricted to making recommendations.

The important question then becomes, how is the information initially available communicated to the capital-allocation mechanisms? There are two classes of reasons why information cannot be conveyed without cost: (1) communication channels have limited capacity; and (2) there are incentive effects that reduce the reliability of information transmission.

1. The specialists who have the concepts and undertake the research have more knowledge of the context than others. An engineer has had training that may not be available to the generalist who allocates resources. Thus, any information conveyed will not be understood as well by the recipient as by the sender. Further, the specialists have spent more time with the project than could any reviewing agency. The agency has too many other responsibilities. The capacity to absorb information is always limited. Hence, again there is a degradation of information with transmission.

   There may, to be sure, be situations in which the central mechanism has better information in some respects than the specialists. It might have better knowledge of other similar concepts and might well have better understanding of the commercial, as opposed to technical, possibilities. However, there will always be a degradation of the technical information, so that the probability distribution of outcomes of the development process (for given development costs) will on the average be wider.
2. Within a given firm it may be assumed, as a first approximation, that there is no distortion of the information; the specialist presents the

information as well as he or she can. However, if information has to cross the boundaries of the firm (for instance, to attract capital from outside investors), the incentive increases to present information misleadingly. Negative aspects might be slurred over, probabilities of success exaggerated.

As a second approximation, there can be some distortion even within the firm. There is some incentive to increase the importance of one's work, to make it appear more valuable in potentiality and thereby earn material and nonmaterial rewards. As in any investment activity, the individual bears limited financial responsibility for failure. Furthermore, for research and development over extended periods of time, the feedback is so slow that the individual is not apt to be in the same position when the program shows results. Finally, the responsibility for success and failure in any position, but especially in one involving such uncertainties as those of research and development, is very hard to assess. That a project failed by no means proves that it should not have been undertaken.

From these considerations, the following implications may be drawn:

1. When responsibility for decisions on research and development is shared because of a need for approval of capital expenditures, the information used in making these decisions is apt to be degraded from its initial state.
2. The longer the chain of communication involved in the approval of projects, the more the information is apt to be degraded.
3. When the chain of communication crosses the boundary of the firm, the degradation of information is apt to be much more severe.[11]

IV. THE DECISION TO INNOVATE AND FIRM SIZE

We can now draw the threads of the analysis together. In particular, the different strategic responses of small firms and large firms to the emergence of research concepts will be analyzed. For simplicity, I speak as if there were just two discrete sizes of firm; of course, there is a continuum of firm sizes and a parallel continuum of innovation strategies.

Innovation has been described as a two-stage decision process. As usual, the appropriate analysis must proceed in reverse order of time. That is, we must first study the decision on development expenditure given the research concept and the research outcome, then analyze the decision to engage in research.

Suppose, then, that we compare a large and a small firm, both of which engaged in research starting from the same concept and observed the same research outcome. The small firm is well informed about the development possibility function; that is, the function relating expected profitability to a given level of development expenditures. It can therefore calculate an optimal level of development expenditures. However, if the amount is large enough, it will not be able to finance it from its own capital funds. It could seek capital from outside. Assume, however, that it has fully utilized whatever general borrowing power it has. Then it has to seek financing based on the project itself. However, for reasons adduced in the last section, the transmission of information across the boundary of the firm will be accompanied by considerable degradation. It follows that capital will be available from the outside only on unfavorable terms (if at all) so that the scale of development expenditures will be less than optimal. Indeed, if the amount of development funds required is very large, it will be essentially impossible to finance the project by borrowing.

A small firm can in many cases obtain outside financing by sale of equity. When the amounts involved are large relative to the initial size of the firm, the transaction amounts to selling the development prospects and is likely to be accompanied by a change in control. (I will take up the possibility of sale of research findings in the next section.)

A large firm facing the same research concept and research outcome will have much less severe restrictions on funding. However, as usual in economic affairs, there is no pure gain without offsets. The difficulties of communication with an external capital market are replaced by those of communicating with the internal capital-allocation mechanism. As we have seen, the information loss in the large firm is greater than that in the small firm, but less than that involved in reaching the external market. Therefore, the larger firm will tend to invest suboptimally in development expenditures. However, it will do better than the small firm for large development expenditures that the large firm can finance but the small firm cannot. It will do less well on expenditures small enough for a small firm to finance.

As an additional hypothesis, it might be supposed that the information loss in the large firm is greater for proposals with greater novelty. The prior information of the internal capital-allocation mechanism may not equip it to evaluate novelties very well. The smaller firm, having less information loss, may be able to accept greater novelty more easily (provided it can finance the development process). Hence, there may be a bias against greater originality in large firms.

It may be objected that a large firm is not more capable of financing large expenditures than a small one. It has larger resources, but it also has

larger demands of all kinds. Hence, it is no more capable of financing a given large expenditure for development than the small firm, as it has other large development expenditures competing for the scarce funds. This is an important point. But there are at least two reasons why we would expect the financing ability of large firms to grow more than proportionately to their size: (1) as we have already seen, large firms have disproportionate access to the external capital market without reference to specific projects. Hence, the pool of available capital is more than proportionately larger. Further (as also noted) the size of the financing available is likely to be statistically steadier, decreasing the probability that a demand for a large amount of development expenditures will coincide with a transient shortage of capital funds. (2) If there are a number of potential demands for development expenditures, the demand will also be statistically steadier. There is a high probability that an above-average demand for development expenditures in one area will be offset by a below- average demand in another. This potential offsetting is less available in small firms.

Basically, then, the superiority of large firms in financing rests on the operations of the insurance principle, though it is aided by economies of information to companies that supply capital to the large firms.

We have first analyzed in dynamic programming form the effects of firm size on the development decision, given the outcome of the research phase. From the above reasoning, for each research concept and each research outcome there is an expenditure on development and a probability distribution of profitabilities in production. These will be affected (as indicated) both by the development profitability function and by the availability of capital, which is conditioned by the problems of information transmission. Hence, there will be a probability distribution of anticipated profitabilities taking account of both development expenditures and subsequent profitabilities in production. It has been argued that, on the average, small firms will be superior if the optimal development costs are low and large firms will be superior if costs are large.

Now consider the decision to engage in research. (Again, for the time being, ignore the possibility that the research outcome can be sold.) Before engaging in research, the development profitability function is itself unknown; nevertheless there will be expectations of it. In probability language, the development profitability function (itself a random variable expressing the distribution of profitability in production conditional on development expenditures) is taken as conditional on research outcome after that is known and as unconditional (more precisely, conditional on research concept but not on outcome) before research is undertaken.

Given the research concept, it may be expected (though without certainty) that subsequent development expenditures will be low if the

project is at all feasible. In that case, it follows that small firms will be more likely to undertake the research than large firms. The opposite is the case if the unconditional distribution implies that development expenditures are likely to be high. Already at the point where the decisions to undertake research are made, there is differential selection among firms of differing sizes.

Thus, on the average one would expect firms to specialize in projects whose optimal development scales are correlated with the size of the firm. Projects anticipated to lead to large expenditures will on the whole be less than optimally funded, because large firms have higher transmission losses for information.

If the supplementary hypothesis advanced above is correct (that larger firms will find it harder to allocate capital to very novel ventures), then it is also true that very novel research concepts will be less likely to lead to research projects in large firms than in small.

Finally, it must be pointed out that the correlation in research undertaken between firm size and optimal level of development expenditures, though positive, will be far from perfect. The level of development expenditures, as repeatedly emphasized, will depend on the research outcome. Research, by its nature, is uncertain. It can easily happen that a research program is undertaken with a probability distribution of optimal development expenditures whose expectation is relatively small before the research outcome is observed. The distribution conditional on research outcome may be quite different, possibly with a large expectation. This is a far from rare event. Of course, the opposite can also occur. If the correct distributions are held, it must occur comparably frequently.

It can therefore happen that a small firm undertakes a line of research whose outcome would optimally involve a much larger development expenditure than it is prepared to undertake. It will either pursue the development on a much smaller scale than optimal, or it will discontinue it altogether if there are sufficient increasing returns to scale in the development process.

V. A MARKET FOR RESEARCH OUTCOMES

This concluding section seeks to remove one limitation of the preceding. The research outcome may itself be the object of a market transaction. Selling ideas is not entirely as simple as selling goods, but they are valuable to at least some potential buyers. Establishment and transfer of properly rights can take several forms. The research outcome might be patentable, in which case the sale is straightforward. Alternatively, the buyer might value a whole constellation of working knowledge embodied in the firm. In that case, the

sale of the research outcome could be equivalent to the sale of the whole firm.

From the discussion thus far, the natural sellers of research outcomes would be small firms that, after observing the outcome, determine that optimal expenditures on development exceed the financial capacity of the firm. In view of the uncertainty about development costs at the moment of the research concept, such situations can arise easily. The buyers might be individuals or groups of individuals in the external capital market who wish to secure their investment in such an uncertain situation by equity acquisition rather than bonds. More likely, however, it is the large firms in similar fields who constitute the natural demand side of the innovation market, whether research outcomes are sold in the form of patents or of whole firms.

The existence of markets for research outcomes alters the incentive structures for undertaking research within both large and small firms. For small firms, it lessens the inhibition on starting research for which large development expenditures are likely. If this came to pass, they do not find the research useless; they can sell the outcome to a large firm. One must still reckon with a loss of information as it passes across the boundaries between the large and small firms. Hence, the incentives for the research are less than they would be within the large firm. Since the large firm is well informed, it is also true that the loss of information is less than it would be between the small firm and the general external-capital market, so that the possibility of sale to large firms is not negligible, as we have assumed the external financing of expensive developments by small firms to be.

The existence of markets for research outcomes also alters the incentives for research within large firms - for the worse. For now the firm has an alternative supply of research outcomes on which to base its development of innovations. The constraints on its total development expenditures imply that anticipated availability of research outcomes on the market will reduce the incentive to use only internally generated research outcomes. There are limits to relying on the market for research inputs into the development process. For example, internal research capability is complementary to externally purchased research outcomes. It is needed to evaluate them and to synthesize them with other research outcomes, whether internal or external. But clearly some substitution takes place.

If this analysis is meaningful, it suggests a division of labor according to firm size. Small firms will tend to specialize more in the research phase and in smaller development processes. Larger firms will devote a much smaller proportion of their research and development budget to the research phase. They will specialize in the larger developments and will buy a considerable fraction of the research basis for their subsequent development of innovations.

While anecdotes are no substitute for good statistical analysis, a striking number of innovations have been produced by giant corporations on the basis of ideas (and perhaps some production) from small firms.

NOTES

1. J.A. Schumpeter, *Business Cycles,* Vol. 1 (New York and London: McGraw-Hill, 1939), pp. 94-109.
2. A.A. Berle, Jr., and G.C. Means, *The Modern Corporation and Private Property* (New York: Macmillan, 1932).
3. K.J. Arrow, and F.H. Hahn, *General Competitive Analysis* (San Francisco and Edinburgh: Holden-Day and Oliver & Boyd, 1971).
4. E.H. Chamberlin, *The Theory of Monopolistic Competition,* 6th ed. (Cambridge, Mass.: Harvard University Press, 1950).
5. See, for example, Table 3.1, p. 40, in F.M. Scherer, *Industrial Market Structure and Economic Performance* (Chicago: Rand McNally, 1970).
6. For the period from 1899 to 1939, see G.W. Nutter, *The Extent of Enterprise Monopoly in the United States 1899-1939* (Chicago: University of Chicago Press, 1951). For more recent trends or lack thereof, see W.F. Mueller and L.C. Hamm, "Trends in Industrial Market Concentration 1947 to 1970." *Review of Economics and Statistics 56* (1974):511-520.
7. G.J. Stigler, "The Division of Labor is Limited by the Extent of the Market," *Journal of Political Economy 56* (1951):185-193.
8. A. Chandler, Jr., *The Visible Hand: The Managerial Revolution in American Business* (Cambridge, Mass: Harvard University Press, 1977).
9. For more complete discussion, see K.J. Arrow, *The Limits of Organization* (New York: Norton, 1974): Ch. 2.
10. R.H. Coase, "The Theory of the Firm." *Economica N.S. 4*(1937):368-405.
11. In order not to interrupt the main line of the argument, I have left rather vague the concept of the profitability of the innovation, which appears as an output of the development process. It is not necessary that the innovation give rise to market power—that is, that it be a commodity with some distinct differentiation from others and on which, therefore, a monopoly profit can be earned. (Of course, this possibility is not excluded either.) But even if the product, or a close substitute, is one already produced, an innovation may amount to a cost reduction. Hence, the firm will earn a rent on the superior productivity induced by the innovation. It must be recognized, however, that the knowledge embodied in an innovation cannot fully be made property. It is apt to be copied by others, and, as the knowledge spreads, the price of the product will decline. Hence, the anticipated profitability must take account of the declining rent from the innovation.