



Local Search and the Evolution of Technological Capabilities

Toby E. Stuart; Joel M. Podolny

Strategic Management Journal, Vol. 17, Special Issue: Evolutionary Perspectives on Strategy (Summer, 1996), 21-38.

Stable URL:

<http://links.jstor.org/sici?sici=0143-2095%28199622%2917%3C21%3ALSATEO%3E2.0.CO%3B2-L>

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

Strategic Management Journal is published by John Wiley & Sons. Please contact the publisher for further permissions regarding the use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/jwiley.html>.

Strategic Management Journal
©1996 John Wiley & Sons, Ltd.

JSTOR and the JSTOR logo are trademarks of JSTOR, and are Registered in the U.S. Patent and Trademark Office. For more information on JSTOR contact jstor-info@umich.edu.

©2003 JSTOR

<http://www.jstor.org/>
Fri Sep 19 10:21:09 2003

LOCAL SEARCH AND THE EVOLUTION OF TECHNOLOGICAL CAPABILITIES

TOBY E. STUART

Graduate School of Business, The University of Chicago, Chicago, Illinois, U.S.A.

JOEL M. PODOLNY

Graduate School of Business, Stanford University, Stanford, California, U.S.A.

The assumption that 'local search' constrains the direction of corporate R&D is central in evolutionary perspectives on technological change and competition. In this paper, we propose a network-analytic approach for identifying the evolution of firms' technological positions. The approach (1) permits graphical and quantitative assessments of the extent to which firms' search behavior is locally bounded, and (2) enables firms to be positioned and grouped according to the similarities in their innovative capabilities. The utility of the proposed framework is demonstrated by an analysis of strategic partnering and the evolution of the technological positions of the 10 largest Japanese semiconductor producers from 1982 to 1992.

INTRODUCTION

A common assumption of evolutionary perspectives on industrial innovation is that 'local search' significantly constrains the direction of corporate R&D (Nelson and Winter, 1973; Dosi, 1988; Teece, 1988; Cohen and Levinthal, 1989). The characterization of search as 'local' or 'problemistic' (Cyert and March, 1963) implies that organizations initiate new R&D projects that share technological content with the outcomes of their prior searches. It seems uncontroversial to assert that the notion of 'local search' is relative: the term *local* presumes a broader context of inventive activity forming the backdrop against which the search behavior of a focal firm can be referenced. However, while the qualifier *local* has meaning only when it is paired with the specification of a broader search context, the literature has yet to provide a generalizable approach for characterizing this technological landscape and the positions of firms within it.

In this paper, we propose a network-analytic

methodology to measure the technological landscape that is produced by the simultaneous search activities of a group of high-technology firms. A firm's position or niche in this landscape derives from the overlap of its inventive activities with those of its competitors. In our approach, a firm engages in search when its niche shifts across time periods, and the manifestation of 'localness' is equivalent to the amount of its niche shift. We propose this approach because it enables a systematic assessment of the extent of interfirm, intertemporal, or interindustry differences in the 'localness' of search.

Our primary objective is to illustrate the methodology's capacity to describe changes in firms' technological positions. However, the approach we present is relevant to other areas of research, particularly theories of the resource-based view of the firm and of strategic groups. In our analysis, firms' technological positions derive from one competence that partly shapes their competitive success: the ability to innovate in particular technological subfields. Specifically, we propose a relational construction of technological positions such that firms that have developed portfolios consisting of similar technologies are located near to one another. Assuming that firms' abilities to

Key words: local search, evolutionary theory, technological change, innovation

develop technologically similar inventions reveal proximities in their underlying 'innovative capabilities,' then firms' technological positions in this paper reflect their innovative capabilities. Moreover, clusters of firms with adjacent technological positions cohere because their members have similar innovative capabilities, and so they can be seen as strategic groups. In the discussion, we suggest that sociologists' notion of a role lends to our construction of technological positions a theoretical basis and represents a compelling approach to measuring clusters of firms. The benefit of this approach is that sociologists have developed well-established techniques for measuring role equivalencies in a network.

The paper is organized according to the following plan. The first section discusses the literature that elaborates diverse organizational causes of local search. The next section develops the methodology for representing local search within a broader context. The third section discusses a data source, patent citations, which is used to measure firms' technological niches and niche shifts. The fourth section introduces the empirical setting and the sample—the Japanese semiconductor industry during a 15-year period. The fifth section contains the maps of technological positions of the sample members, and it relates position in these maps to the market shares, number of patents, and a measure of the innovativeness of the sampled firms. The sixth section is a discussion that draws parallels between the approach developed in this paper and the resource-based view of the firm. The final section elaborates implications of and extensions to this research.

LOCAL SEARCH IN R&D

The literatures on evolutionary economics, the management of technology and organizational theory, all posit that R&D is history dependent. In other words, organizations search for novel technologies in areas that enable them to build upon their established technological base. This local search results from individual and organizational level processes, as well as from the nature of the firm's innovative capabilities.

At the level of the individual decision maker, bounded rationality engenders local search when organizational members fail to consider the uni-

verse of possible applications of R&D funds and, instead, look to the firm's previous development decisions for guidance. The management of R&D involves investment decisions that must be made in the context of uncertain technical, economic, and social environments in which the actions of competitors are particularly difficult to anticipate (MacKenzie, 1992; Tushman and Rosenkopf, 1992). In such ambiguous and uncertain settings, a heavy reliance on historical experience is the norm (March, 1988). In other words, the results of past searches become natural starting points for initiating new searches (Nelson and Winter, 1982).

At the organizational level, local search is produced by the smooth functioning of organizational routines (Cyert and March, 1963; Nelson and Winter, 1982). A routine is defined by Nelson and Winter (1982: 96) to be a pattern of activity that is repeatedly invoked. In Nelson and Winter's schema, routines generate similar organizational responses to frequently encountered stimuli, and are therefore the source of continuity in organizational behaviors. The upshot of this conception of the firm is that organizational behaviors like R&D are delimited by the routines that evolve in a firm. Even when environmental conditions have decreased the attractiveness of a particular activity to a firm in possession of a given skill set, intraorganizational politics and historical precedent can prevent or slow managers from abandoning a particular technical undertaking (Burgelman, 1994).

Another reason that search is likely to be local is that organizations have a higher likelihood of successful technology development in areas in which they have prior experience. Organizational learning is a cumulative activity that is facilitated by concentrating it in areas of prior knowledge accumulation. The competence to innovate in a particular domain follows consistent investments to develop the facilities, personnel, intellectual property, interorganizational relations, and tacit organizational knowledge to successfully innovate in that technological area (Teece, 1988). This means that the knowledge stock a firm has accumulated in a technological subfield conditions its returns to R&D investments in that subfield (Cohen and Levinthal, 1989). Therefore, it is natural to expect that R&D will produce superior results when it is concentrated in the areas of a firm's established competencies.

Historical, case study, and other empirical research provide scattered evidence to support the hypothesis of local search in many technological areas. Even when a major shift in technology strategy is desired, the literature proposes a number of reasons why firms may have a limited ability to make rapid adjustments. Lee and Allen (1982) showed that one firm required a number of years to integrate new technical staff, suggesting that it may take a considerable amount of time for organizations to acquire and assimilate new technological knowledge by augmenting or making substitutions in their staff of technologists. There is also evidence to show that high-tech firms do not capriciously shift the market niches in which they participate. In a study of the semiconductor industry, Boeker (1989) found that entrepreneurial firms typically maintained the strategies that they had at the time of founding. Podolny and Stuart (1995) found that semiconductor technologies in crowded technological areas were the ones most likely to be elaborated in later periods because they were within reach of the search areas of many firms.

The constraint of local search is also implied by conceptual frameworks that have highlighted the difficulties experienced by incumbent firms in adjusting their technology strategies to major environmental changes (Abernathy and Clark, 1985; Tushman and Anderson, 1986; Henderson and Clark, 1990). These studies have discussed and documented the effects of 'competence destroying' technical changes, which are defined as major technological changes that obviate the technical competencies of established firms. A finding of these studies is that when radical technological developments shift the basis of competition, the path-dependent nature of firms' capabilities prevents them from responding quickly. Importantly, such observations do not suggest that there is no variation in a firm's technological developments, but they strongly imply that a firm's technical developments do not follow sudden and unanticipated changes.

This review of the literature has been devoted to establishing the widespread prevalence of the assumption of local search. Nevertheless, it remains the case that the primary empirical evidence to support this assumption comes from in-depth case studies of individual organizations or industries (Abernathy and Clark, 1985; Burgelman, 1994; Helfat, 1994; Rosenberg, 1969; Sahal,

1985). Qualitative studies with the firm as the unit of analysis (e.g., Burgelman, 1994) have documented the history-dependent quality of corporate R&D, while those concentrating on the industry or technical field (e.g., Sahal, 1985) have traced the path-dependent nature of industry- or field-level technical change. Although these studies richly describe organizational learning and technological evolution in specific historical periods, the methodologies that they employ do not lend themselves to a systematic assessment of interfirm, intertemporal, or interindustry variance in the scope of search.¹ Without a generalizable method allowing for such a systematic assessment, it is difficult to (1) identify which members of a group of competitors have been the most locally bounded in the outcomes of their R&D, (2) measure the extent to which the search trajectories of the members of a group of firms converge or diverge over time, or (3) test basic hypotheses of how a firm's technological position at one point in time is contingent on its prior position and search trajectory.

NICHE OVERLAP AND EVOLVING TECHNOLOGICAL POSITIONS

Clearly, the appropriate place to look to assess the degree of path dependence in corporate innovation is the actual technological knowledge created by a firm. In this section, we develop a methodology in which all of the recent inventions of a group of firms serve as a reference point for identifying relative technological shifts of individual members. We propose that companies which shift technological positions relative to their competitors are the ones that have moved the greatest technological distance from the po-

¹ An alternative approach has been to explore the implications of local search in simulation studies (Nelson and Winter, 1982; Winter, 1984). Simulations have typically situated corporate search in the context of an abstract space identified by standard economic variables, such as input coefficient magnitudes (Winter, 1984). This approach entails defining the context of search as a probability distribution that represents a set of input coefficients *in the neighborhood of a firm's current production techniques*. In other words, the terrain over which search takes place is an 'economic space' of input coefficients, and the degree to which localness is built into the model is reflected in the parameters of the search distribution. Given the assumptions of this approach, it is simple to assess interfirm distances and the rate and direction of a firm's movement in 'economic space'.

sitions that they had previously occupied. These are the companies that have deviated from locally circumscribed research.

Our methodology allows each firm to occupy a ‘technological niche’ that emerges from the distribution of technological antecedents of the firm’s current technology developments (Stuart, 1995; Podolny, Stuart and Hannan, 1996). We define the technological overlap between the members of a pair of firms in terms of the extent to which they build on the same foundations for their current inventions. We will use the notation α_{ij} to denote the proportion of firm i ’s niche that is occupied by another firm j : α_{ij} represents the proportion of inventions built upon by firm i that are also foundations for the inventions of j . Therefore, α_{ij} is bounded by zero and one: at zero, two firms are completely differentiated; at one, j fully occupies i ’s niche.

For a system of N innovators, complete information about interfirm technological overlaps can be expressed in an asymmetric matrix of order $N \times N$ (McPherson, 1983; Hannan and Freeman, 1989). The elements of this matrix are called ‘competition coefficients,’ and the matrix itself is known in the literature as a ‘community matrix.’ The competition coefficients are simply the α_{ij} , α_{ji} for ($j = 1, 2, \dots, N; i = 1, 2, \dots, N; i \neq j$).

Figure 1 depicts a hypothetical technological network including three firms, denoted A, B, and C. The figure includes arrows, which represent technological building relations at the level of the discrete invention. The arrows are directed from the firms to a number of inventions that belong to unidentified actors; each arrow represents the act of building on a discrete invention. For example, four inventions were foundations for A’s technologies. In the hypothetical network, α_{AB} is 0.5 because B builds on two of the four inventions that are foundations for the technologies of A.

Figure 1 also illustrates the corresponding community matrix for the three networked firms. The first row of this matrix registers the degree to which each of the companies in the sample occupies the niche of firm A. Thus, B occupies 50 percent and C occupies 75 percent of A’s niche. The first column indicates the extent to which firm A occupies the niches of the other members of its network. Thus, A overlaps with 100 percent of B’s and 37.5 percent of C’s niche. The main diagonal has no significance and so it is set to missing.

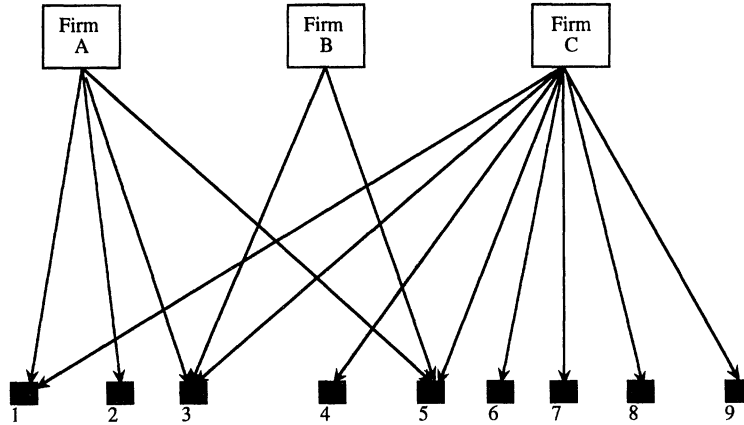
The community matrix will be denoted \mathbf{A}_{t_m} . The measure of overlap that we use produces asymmetric competition coefficients in each pair of firms. The elements of the \mathbf{A}_{t_m} matrix for the ij th dyad at time t_m are defined to be:

$$\alpha_{ijt_m} = \frac{\sum_{v=1}^p a_{ivt_m} a_{jvt_m}}{\sum_{v=1}^p a_{ivt_m}} \quad (1)$$

$$\alpha_{jit_m} = \frac{\sum_{v=1}^p a_{ivt_m} a_{jvt_m}}{\sum_{v=1}^p a_{jvt_m}} \quad (2)$$

where v denotes a technological antecedent, and p indexes the total number of distinct antecedents that were foundations for the sampled firms at time t_m . The value a_{ivt_m} is coded 1 if antecedent v served as a foundation for the inventions of firm i at time t_m , and 0 otherwise; similarly, a_{jvt_m} is coded 1 if antecedent v is a foundation for the inventive activity of firm i at time and 0 otherwise. Two firms, i and j , produce both an ij and a ji cell in the \mathbf{A}_{t_m} matrix. The ij th cell results from counting the number of common antecedents of i and j ’s inventions at time t_m and then dividing this sum by the total number of distinct technological precursors of firm i ’s activity. Similarly, the ji th element results from taking the same numerator, but in this case dividing by the total number of antecedents of j ’s activity. Clearly, the ij th and ji th cells in the \mathbf{A}_{t_m} matrix generally will not be equal to one another; although the numerator is common to both cells, the denominator in almost all cases will differ (one exception is if there is no overlap in the antecedents of i and j , in which case both cells equal zero).

Given that the ij th cell represents the degree to which firm j is in firm i ’s niche, it should be clear that row i specifies the extent to which all other firms are in i ’s niche, and column i specifies the degree to which firm i is present in the niches of all other firms. (Returning to the hypothetical community matrix of Figure 1, row 1 specifies the occupants of firm A’s niche and column 1 registers A’s presence in the niches of its alters). Taken together, row i and column i define the



Corresponding Community Matrix:

Firm	A	B	C
A	0	.50	.75
B	1	0	1
C	.375	.25	0

Figure 1. Hypothetical technological network for three firms. The figure illustrates the level of competitive crowding among three hypothetical firms, A, B, and C. In the figure, the objects of the arrows emanating from the firms, the numbered boxes, represent existing inventions. The lines with arrows represent technological building relations. For example, firm B has developed inventions that built on inventions 3 and 5. Firm A's row in the community matrix registers the percentage of its niche filled by B and C. Firm A's column in the matrix is the percentage of the niches of B and C that it occupies

technological position of a focal firm with respect to all other firms at a particular time t_m . In effect, the entries in row and column i define a *global position* for firm i in a $2N - 2$ dimensional space, where N is the number of firms in the community matrix.

This conception of a firm's global position as a function of the proximities of its technological developments to those of the members of a group of competing firms serves as our point of departure for measuring the technological distance between firms in each period of time. In addition, we will use this measure of position to define the extent of a firm's technological movement across time periods. Specifically, we define the distance between i and j at time t_m in terms of the degree to which i and j have a similar pattern of niche overlap with all other firms k . Formally, the (Euclidean) distance between firm i and j for a given time t_m is defined to be:

$$d_{jit_m} \equiv d_{ijt_m} = \left\{ \sum_{k=1}^n [(\alpha_{ikt_m} - \alpha_{jkt_m})^2 + (\alpha_{kit_m} - \alpha_{kjt_m})^2] \right\}^{1/2}, \quad k \neq i, j \quad (3)$$

where the alphas are the (asymmetric) competition coefficients for the ik th and jk th dyads at time t_m . Notice that the distance between firms i and j in Equation 3 is a function of the level of the dissimilarity of their patterns of niche overlap with each of the other $(N - 2)$ firms in the sample. Thus, $(\alpha_{kit_m} - \alpha_{kjt_m})$ is the difference in the extent to which firms i and j occupy the niche of a third firm k , and $(\alpha_{ikt_m} - \alpha_{jkt_m})$ is the difference in the extent to which the niche of k overlaps the niches of i and j .

Similarly, it is possible to quantify the intertemporal shift of firm i 's technological niche in terms of the degree to which its pattern of niche overlap changes over time. Formally, we define the shift in firm i 's technological niche from time t_l to t_m as:

$$d_{it_it_m} \equiv d_{it_mit_l} = \left(\frac{n-2}{n-1} \right) \left\{ \sum_{k=1}^n [(\alpha_{ikt_l} - \alpha_{ikt_m})^2 + (\alpha_{kit_l} - \alpha_{kit_m})^2] \right\}^{1/2}, k \neq i \quad (4)$$

The expression inside of the sum operator registers the extent to which firm i 's niche overlap with the other $(n-1)$ firms changes between time periods t_l and t_m . The more that the pattern of i 's overlap changes between t_l and t_m , the larger will be the summed expression of Equation 4.²

Equation 3 represents the distance between different firms within a single time period. Equation 4 yields the amount of a single firm's niche shift across time periods (i.e., the distance between the positions occupied by firm i in period t_l and the same firm in period t_m). Finally, to represent the distance between different firms in different time periods, we construct a symmetric matrix, \mathbf{D} , where cell it_jt_m registers the difference in the pattern of overlap between firm i at time t_l and firm j at time t_m . Formally, we define the elements of the matrix \mathbf{D} :

$$d_{it_jt_m} \equiv d_{j_t_mit_l} = \left(\frac{n-2}{n-1} \right)^\delta \left\{ \sum_{k=1}^n [(\alpha_{ikt_l} - \alpha_{jkt_m})^2 + (\alpha_{kit_l} - \alpha_{kjt_m})^2] \right\}^{1/2}, k \neq i, j \quad (5)$$

where δ equals 1 if $i = j$ and $l \neq m$, and 0 otherwise. According to Equation 5, the more that firm i 's pattern of niche overlap with its competitors in period t_m is similar to firm j 's pattern of overlap with its competitors in period

t_l , the lower will be the value of cells $d_{it_jt_m}$, $d_{j_t_mit_l}$.

Equation 5 incorporates the specifications of Equations 3 and 4. Specifically, when $\delta = 1$, $t_l \neq t_m$ and $i = j$, Equation 5 reduces to Equation 4, and when $\delta = 0$ and $t_l = t_m$, Equation 5 reduces to Equation 3. Assuming that all firms are present for all time periods, the dimensions of the symmetric \mathbf{D} matrix are $N * T$ rows by $N * T$ columns, where N is the number of firms and T is the number of time periods. Given the nested equations, Equation 5 identifies a matrix that includes three types of information: (i) the distance between all firms within time periods, (ii) the distance between each firm and *itself* across time periods, and (iii) the distance between different firms across different time periods.

Readers familiar with the social network literature will recognize the Euclidean distances of Equations 3, 4 and 5 as continuous measures of structural equivalence. Structural equivalence is a measure of the extent to which two actors are closely situated in their network because they have similar ties to the other network members. As Burt (1987) observed, the more similar are the relational patterns of two network members, the greater is their structural equivalence and therefore the more that one member could substitute for the other member in its role relations. In effect, our approach defines the context of a focal firm's search by the technological undertakings of competing firms. Search can be considered to be a structural property in that a focal firm's change in position across periods of time can be defined by its niche shift between times t_l and t_m .

Using conventional multidimensional scaling (MDS) routines, it is possible to convert the information in the \mathbf{D} matrix to a graphical representation of interfirm distances. However, it is first necessary to construct the competition coefficients (the \mathbf{A}_{t_m} matrices) from which technological distances can be derived. To do this, we use the patent citations made by a sample of semiconductor firms.

PATENTS AND TECHNOLOGICAL LINEAGE

Patents identify inventions because they are only granted to products, processes, or designs that are industrially useful and nonobvious to an individual who is knowledgeable in the relevant technical field. An important component of the patent

² The sum in Equation 4 is multiplied by $((n-2)/(n-1))$ so that the metric is comparable to that in Equation 3. When the Euclidean distance between i and j is measured, the comparison is across $n-2$ other actors. However, when the Euclidean distance between i at time period t_l and i at t_m is assessed, there are $n-1$ comparisons. Since there are more comparisons when i is compared to itself across time period, we deflate the distance by $(n-2)/(n-1)$ so that the distance that a firm shifts over time is comparable to the distances between firms within a particular time.

application procedure is the ‘prior art’ provision. In the United States, previous U.S. patents that are identified as technological precursors to the current invention are referred to as ‘prior art.’ The citation process is legally important because it limits the claims of a pending patent: legal protection is awarded only to the technological claims that are not anticipated by the prior art. A number of scholars have noted that patent citations trace out technological building relationships among inventions (e.g., Jaffe, Trajtenberg and Henderson, 1993).

Given that patent citations identify the technological antecedents of a firm’s current inventions, we use patent citations to quantify technological niche overlaps among a community of innovating firms. Recall that α_{ij} was used to represent the extent to which the inventions of firm j shared antecedents with firm i . Because patent citations identify technological building relations, we measure the α_{ij} as the proportion of patents cited by i that are also cited by j . For example, if i cites 100 patents and j cites 50 of those patents, α_{ij} equals 0.5.

SETTING: THE JAPANESE SEMICONDUCTOR INDUSTRY

To illustrate the utility of this methodology for identifying search trajectories, we map the technological positions of the largest firms in the Japanese semiconductor industry. A number of considerations motivated the choice of this setting. First, the semiconductor industry is one that is still very much technology-driven. Semiconductor production involves tremendously complex processes (Langlois *et al.*, 1988), and technical advances have incessantly driven down the price and increased the performance of semiconductor devices throughout the history of the industry. For this reason, R&D expenditures are quite high (routinely exceeding 10% of revenues for many incumbents), and firms’ decisions about which technological area(s) to target are critical factors in determining organizational performance. Second, the Japanese industry underwent radical change during the period of the analysis. Our data span the period from 1978 to 1992. Although a few Japanese firms began semiconductor production in the 1950s, they were comparatively minor players in the global marketplace until the

late 1970s and early 1980s. Therefore, we observe the evolution of the industry during the interval in which it achieved global prominence. Finally, a number of detailed books on the Japanese industry offer a yardstick against which to compare the results of this analysis.

The sample includes the 10 largest Japanese semiconductor manufacturers: Fujitsu, Hitachi, Matsushita, Mitsubishi, NEC, Oki Electric, Sanyo, Sharp, Sony, and Toshiba. These firms were vertically integrated, and there was substantial overlap among them in their participation in end-use markets. For example, all of these companies produced computers and consumer electronics products, and most had telecommunications operations.

The data for the analysis are the U.S. semiconductor patents held by each of the 10 sampled Japanese producers. The United States is the world’s largest technology marketplace, and for this reason non-U.S.-based firms routinely submit patent applications in the United States. Each of the 10 sampled firms are among the largest U.S. patent holders for semiconductor device, design, and process innovations. The semiconductor patents held by these 10 firms were collected for the period from 1978 to 1992, inclusive.³

Following the preceding discussion, at a time t_m the matrix of competition coefficients for the network formed by the 10 sampled firms is a 10×10 in which each element registers the extent to which the row firm overlaps with the column firm in its patent citations. One measurement issue encountered in computing the A_{t_m} matrices is the length of time during which the competition coefficients specified by Equations 1 and 2 are calculated. It is unreasonable to define an organization’s technological focus at time t_m only by the inventions that it had patented during the previous year. We therefore chose to create the A_{t_m} matrices from the patent citations made by

³ Semiconductor patents were retrieved from the Micropatent CD series. This series contains all patents granted in the United States since 1976. When a patent is granted, the patent examiner assigns it to a primary class and subclass. The patent is also typically cross-referenced in a number of other classes. We identified approximately 2400 patent class/subclass combinations that included semiconductor device, design, or process inventions, and we included in the dataset all semiconductor patents held by the sampled firms that were either primary-classed or cross-referenced in any one of these locations. Details of the dataset and a list of the 2400 classes are available from the first author.

the firms in the sample during five-year, moving windows.⁴ Our analysis spans the period from 1978 to 1992, inclusive. Using a 5-year window allows us to construct three community matrices—each one derived from nonoverlapping years of data—during the 15-year span of our analysis. Thus, we constructed an A_{82} matrix that is a 10×10 computed from all of the U.S. semiconductor patents awarded to the sampled firms during the period from 1978 to 1982, inclusive. Similarly, the A_{92} matrix was generated from the patent cocitations in the sample during the years from 1988 to 1992.

ANALYSIS

Although assessing search trajectories requires that we consider all years simultaneously, we begin by examining each year separately. As Equation 3 specifies, we construct a separate distance matrix, D_{t_m} , for each of the years 1982, 1987, and 1992. The three panels of Figure 2 show the MDS configurations for each of these years. The coordinates for these plots were generated by the MDS procedure in SAS, version 6.09. In all cases, the number of dimensions was set to two, which resulted in reasonably good stress levels.⁵

Evolution of the technological landscape

As anticipated by the arguments about the path-dependent quality of organizational innovation, the figures suggest a significant degree of stability in the relative positions of the firms in the period

of analysis. In all three figures, the leading-edge semiconductor producers are located to the east of the less technically advanced firms. The firms with the greatest percentage of their electronics end-use business concentrated in consumer electronics products are positioned toward the west end of the figures. This is clearest in panel C, which suggests a two-tier structure. The generalists and technological leaders in 1992 were Toshiba, Hitachi, NEC, Fujitsu, and Mitsubishi. In panel C of Figure 2, these firms occupy positions that appear to be differentiated from their consumer electronics-oriented competitors.

Although Figure 2 suggests a great degree of stability in the structure of the industry, visual comparisons of the panels of the figure are difficult because absolute locations in each of the panels are meaningless and the range of the axes of the panels differ. Therefore, the apparent interfirm distances are not constant across the three panels. To make intertemporal comparisons, we apply MDS to a pooled distance matrix as specified by Equation 5.

Quantifying the amount of niche shift

The result of the MDS of the pooled distance matrix is shown in Figure 3, which confirms that the pattern of niche overlap in the Japanese semiconductor industry was indeed quite stable. Evidence of stability comes from the fact that the position of a firm at time t_m is generally quite close to its position at times t_{m-5} or t_{m+5} . For example, Sanyo in 1987 is relatively near to Sanyo in 1982 and Sanyo in 1992. It is possible to quantify the amount of movement in a firm's position by assessing the change in its column (or row) of the distance matrices specified by Equation 3 at different points in time. One such measure follows this reasoning: if firm i did not change positions over time, then its distance from other firms will be relatively stable (the i th column in the distance matrices at two points in time will be highly correlated).

Following Burt (1988), we assess the stability of a firm's position across time periods by constructing a firm-specific covariance matrix in which each cell represents the covariance between a firm's vector of distances to its competitors across two time periods (therefore, for 10 firms and three time periods, we construct a total of 10 3×3 covariance matrices). The more similar

⁴ We selected 5 years because it is roughly the duration of the product life cycle in the semiconductor industry. For many types of products, five years understates the time interval in which the product is manufactured (e.g., each successive generation of computer memory, 64K, 256K, etc., has been in production for about a decade). However, 5 years may overstate the time period during which a particular product, design, or process is on the leading edge of the technology in the industry (e.g., the next generation of computer memory chip has arrived approximately every 2.5 years).

⁵ In MDS, stress is a normalized, residual sum of squares that suggests the degree to which the resultant configuration agrees with the n -dimensional distance matrix. Stress is often known as a 'badness-of-fit' criterion because higher values suggest worse fits. For the figure representing 1982, the badness-of-fit criterion was 0.086 for two dimensions; for 1987, the badness-of-fit criterion was 0.092; and for 1992, the stress level was 0.036. These stress levels are considered fairly good (Kruskal, 1964).




Figure 2. Technological positions of Japanese semiconductor firms

are a firm's distance vectors across all three time periods, the higher is the percentage of variance captured by the first factor of a principal components analysis of each 3×3 firm-specific matrix. The results of this analysis are reported in Table 1. The findings show that Toshiba experienced the least movement: 92 percent of the variance in its position vectors for the 3 years is captured by the first principal component. On the other hand, Mitsubishi and Fujitsu were the companies that moved the most: the first principal component captured about 75 percent of the variance in the

positions of these two firms. Of all of the sampled firms, Mitsubishi is the only one for which one of the three firm years—1992—had a negative loading on the first principal component. In other words, Mitsubishi moved significantly between 1987 and 1992.

For two reasons, it is remarkable that the pattern of interfirm niche overlap remained so stable during the interval of our study. First, the nature of semiconductor technology is such that a semiconductor device generation change is typically accompanied by significant changes in product

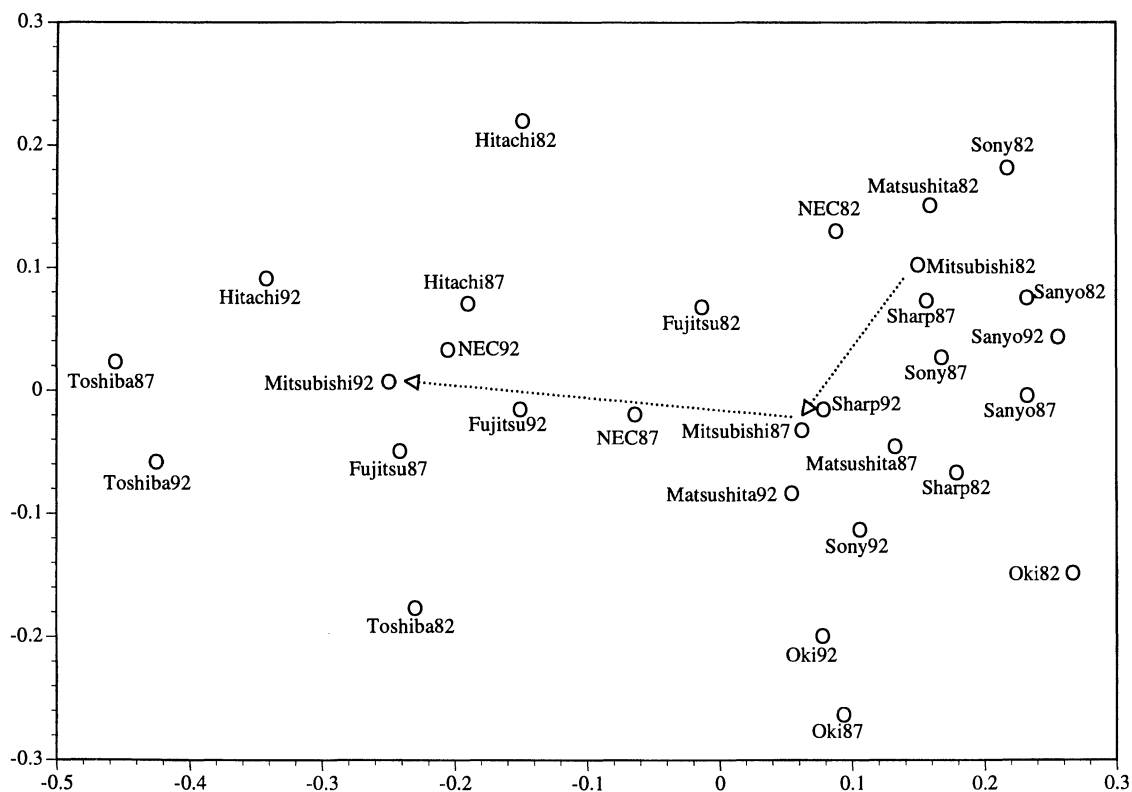


Figure 3. Technological positions of Japanese semiconductor firms: 1982, 1987, and 1992

Table 1. Results from factor analysis to quantify the amount of firms' niche shifts

Firm	Factor 1 ^a	Loading: 1982 ^b	Loading: 1987 ^b	Loading: 1992 ^b
Fujitsu	0.75	0.96	0.41	0.79
Hitachi	0.91	0.91	0.88	0.98
Matsushita	0.80	0.95	0.88	0.84
Mitsubishi	0.75	0.95	0.92	-0.70
NEC	0.78	0.93	0.82	0.89
Oki	0.84	0.92	0.90	0.92
Sanyo	0.90	0.97	0.94	0.94
Sharp	0.85	0.94	0.85	0.93
Sony	0.86	0.95	0.94	0.89
Toshiba	0.92	0.98	0.93	0.95

^aThe first factor indicates the stability of firm *i*'s position. The higher the factor, the more correlated are *i*'s distance vectors across time periods.

^bFactor loadings indicate the extent of a firm's relative movement in a given year.

designs, processes, materials, and manufacturing (it has even been the case that new device generations have required more complex production equipment due to the tighter design rules of more advanced chips). In addition to the fundamental change in the technology from the first year of our data to the last year, we follow the Japanese

semiconductor industry during the period when it grew from a comparatively small size to one of global prominence. During the first year of the analysis, the value of their combined production was under \$2 billion; by the last year, the 10 sampled firms generated \$30.25 billion in semiconductor sales. Despite the dramatic changes in

the scope of the sampled firms and in the industry's technology, the pattern of interfirm technological overlap has remained relatively stable.

Interpreting the configurations

The discussion of movement in different directions in Figure 3 begs the question: What are the implications to a firm of being located in different regions of the technology space? Descriptive accounts of the Japanese industry (e.g., Kimura, 1988; Langlois *et al.*, 1988) help to interpret different neighborhoods of the configurations. The semiconductor operations of Matsushita, Sony, Sharp, and Sanyo were catered to their consumer electronics products businesses. Thus, these companies focused on linear integrated circuits and discrete devices, and so it is not surprising to find that they cluster in one neighborhood of the configurations (see Figures 2 and 3). The technological leaders of the sample were NEC, Hitachi, Toshiba, and Fujitsu. These firms were all broad-line semiconductor producers, but they concentrated on complex devices such as logic circuits and MOS memories to support their operations in computing. In Figure 3, these firms appear to be differentiated from the consumer electronics products companies.

Oki and Mitsubishi are interesting cases because they do not fit neatly with the technological leaders or the consumer electronics products firms. Oki manufactured telecommunications equipment and an array of peripheral equipment for data-processing systems and computers. Therefore, its end-use businesses were close to those of NEC. Nevertheless, Oki possessed neither the breadth nor the level of leading-edge technology of NEC, Hitachi, Toshiba, or Fujitsu. For these reasons, Oki occupied a relatively isolated position in the technological structure of the industry: although the foci of its operations paralleled those of the technology leaders, it played a more peripheral role in the evolution of the industry's technology.⁶

Because Mitsubishi was the firm that sold the greatest percentage of its semiconductor production on the merchant market (according to Kimura, 1988, Mitsubishi consumed only 30% of its semiconductor production in the mid-1980s), its semiconductor focus was not as strongly tied to its production of electronic end-use systems. During the 1970s, Mitsubishi focused on discrete devices and integrated circuits for consumer electronics products. However, following a strategic assessment near the end of the decade, Mitsubishi targeted semiconductors for computer and industrial applications and it augmented its capital and R&D expenditures (Langlois *et al.*, 1988). Around this time, Mitsubishi moved into the DRAM market, developed complementary MOS technology, and began to second source Intel's microprocessors. Figure 3 suggests that the company succeeded in its strategy. Mitsubishi was the single company to exit the group of consumer electronics products firms and join the technological leaders. In Figure 3, the trajectory of Mitsubishi's position shift is highlighted by the arrows that display its movement between each of the time periods.

It is a simple extension of the methodology that we propose to generate 'egocentric' representations of each firm's position. Figure 4 illustrates egocentric perspectives of Mitsubishi's position for each of the three time periods. To generate this figure, we constructed three 9×9 matrices for the years 1982, 1987, and 1992. The 'space' that these matrices represent spans only Mitsubishi's patent citations. In other words, the configurations are representations of how the nine other firms in the sample are distributed through the areas of Mitsubishi's inventive activities in each of these years. In the data matrices for Figure 4, the distance between the firms comprising any particular dyad (e.g., Sharp and Matsushita) is a function of the level of cocitations among those two firms, subject to the limitation that the cocitation must have been of a patent that was also cited by Mitsubishi. Clusters in the panels of Figure 4 represent concentrations of firms that overlap with Mitsubishi's niche in a similar fashion (e.g., they overlap with Mitsubishi in similar technological areas).

The successive panels in Figure 4 illustrate the significant amount of change in Mitsubishi's relative position. In the first panel (1982), NEC, Oki and Sony are isolates: they have no overlap

⁶ For this reason, Oki appears to be less of an isolate in MDSs of a correlation matrix (instead of a Euclidean distance matrix) because correlations eliminate scale effects (i.e., the correlation between the elements of two firms' rows and columns of the community matrix does not reflect differences in their means). However, we believe that it is undesirable to generate the configurations from correlation matrices because scale is an important attribute of firms' positions.

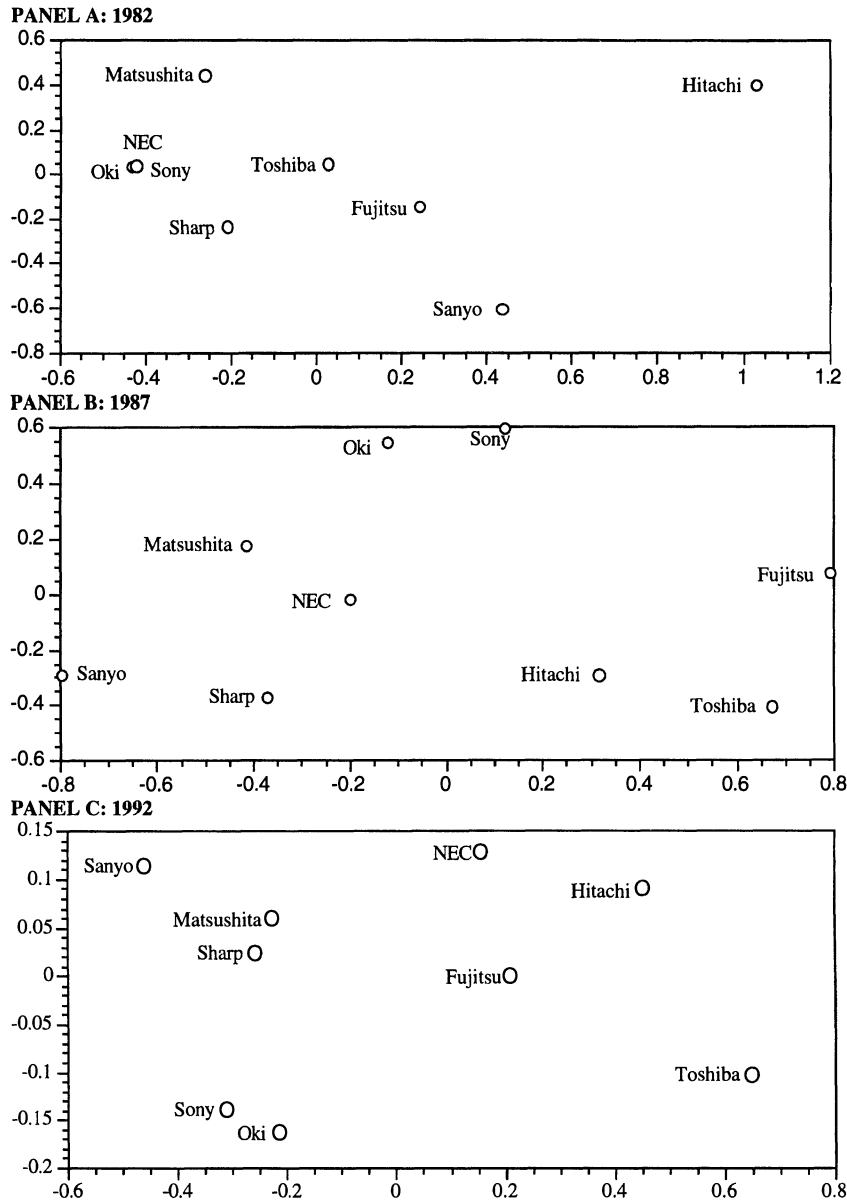


Figure 4. Competitors in Mitsubishi's technological space: 1982, 1987, 1992

at all with Mitsubishi, and so they cluster in the figure. All of the other firms have some overlap with Mitsubishi, but by and large they do not form any discernible pattern based on firm characteristics. By the time of the second panel (1987), all of the nine companies have some overlap with Mitsubishi, so there are no longer any isolates. However, there are still no salient competitive groupings and companies are relatively dispersed in the space. In contrast, in the configuration for 1992 the consumer

electronics/broad-line producer distinction is evident along the east-west axis in panel C of Figure 4. By 1992 the egocentric snapshot of Mitsubishi's position reveals two general groupings of competitors: on one side are the consumer electronics products-focused firms (Sanyo, Sharp, Sony, Oki and Matsushita) and on the other are the broad-line producers (NEC, Hitachi, Toshiba, and Fujitsu).

Returning to the 3-year configuration represented in Figure 3, it is possible to draw axes




Figure 5. 'Regions' in the technological map of the industry

through the figure that associate characteristics of the firms with their positions in the configuration. Specifically, directions in an MDS configuration can be interpreted by regressing variables over the coordinates of the configuration. Figure 5 adds to three regressions over the coordinates; one is for patents, a second is for market share, and a third is for patent citations.⁷ Precisely, the following three regressions were estimated:

$$\text{Share}_{it_m} = \beta_1 \cdot \text{Dim1}_{it_m} + \beta_2 \cdot \text{Dim2}_{it_m}$$

$$\text{Patents}_{it_m} = \beta_1 \cdot \text{Dim1}_{it_m} + \beta_2 \cdot \text{Dim2}_{it_m}$$

⁷ The number of citations received by a patent is a commonly employed measure of the commercial and technical importance of that innovation (Albert *et al.*, 1991). Therefore, the frequency at which a firm's patent portfolio is cited is a combined (i.e., summed) indicator of the importance of its individual inventions. All of these variables (sales, patent cites, and total patents) are measured as proportions to prevent escalation in their values simply as a function of time. Thus, sales are included as market share, patents as the proportion of all patents awarded to a focal firm, and cites as the proportion of all citations that are received by the patent portfolio to a focal firm.

$$\text{Citations}_{it_m} = \beta_1 \cdot \text{Dim1}_{it_m} + \beta_2 \cdot \text{Dim2}_{it_m}$$

where Dim1 and Dim2 are the MDS coordinates of firm i at time t_m . All three regressions had significant F -values and high coefficients of multiple correlation.

The regression analysis corroborates that the firms with the largest number of patents, the highest market share, and the most technologically important inventions are located in the western half of the configuration, angled slightly toward the north. The proportion of all patents received by the firm is the axis with the gentlest slope relative to the horizontal plane; the proportion of patent citations received by a firm is the adjacent axis; and market share is the steepest axis. Bisecting each of the regression lines with a perpendicular divides the configuration into halves. We use this technique to bifurcate the competitive space into two regions, which are quite consistent with the descriptive accounts that distinguish the technological leaders from the consumer electronics products firms. The 'core'

segment is the one that includes the most innovative firms and those with the largest market share. We delineate the two-tiered structure by the bold-faced line in Figure 5 (the market share axis)—this line segments the industry such that the firms with the highest market share are northwest of the bold-faced market share axis.

Strategic positions and interfirm alliances

A number of scholars have suggested that the competitive position occupied by a firm influences its strategic behavior. Specifically in the domain of technology strategy, Kimura (1989) argued that technological position may explain variation across firms in their foreign direct investment activities. Eisenhardt and Schoonhoven (1996) and Shan (1990) hypothesized that the technological position of firms affects their incentives and propensities to engage in interfirm strategic alliances.

We briefly consider the relationship between competitive position and alliance behavior. During the period of the analysis, an exhaustive literature search uncovered 35 alliances involving some type of technology exchange among the semiconductor operations of the firms in the sample.⁸ In Figure 6 we illustrate the pattern of alliances as it relates to the technological positions of the sampled firms. In the figure, the positions of two firms in 1982 were connected with a line if they formed an alliance during the period from 1982 to 1986 (e.g., NEC82 and Oki82). Similarly, two firms in 1987 were linked if they established an alliance between 1987 and 1991 (e.g., Toshiba87 and Hitachi87). Companies that formed a partnership in 1992 were connected for that year, the last year for which we possess this data. Bold lines join firms that engaged in two or more alliances during a time period.

A number of findings emerge from Figure 6. First, it is remarkable the degree to which the 'core' firms—NEC, Fujitsu, Hitachi, and Toshiba—are central in the alliance network. In total, 31 of the 35 partnerships involved one (or two) of those firms. In other words, the pattern of ties is nearly exclusively core-to-core or core-to-periphery. Moreover, each of the four alliances

among the noncore firms included Mitsubishi in 1987. This was the year just prior to the time that Mitsubishi moved into the region of the configuration occupied by the core producers. A related observation about the pattern of intercorporate alliances is that NEC in 1982 and Mitsubishi in 1987, two firm-years that were near the core-periphery border, were particularly active participants in the alliance network. These two firms participated in the greatest number of partnerships among all of the sampled firms in all three years.

Clearly, there is a relationship between position in the configuration of Figure 6 and the decision of a firm to participate in the alliance network. In addition to the fact that alliances appear to bridge the core-periphery border or to join core firms, change in position has a clear relationship to active participation in the recorded technology-exchange and technology-development alliances. For the first period, the correlation between the number of alliances formed by a company and the amount it moved from 1982 to 1987 is 0.13 (not significant). However, for the period from 1987 to 1992, this correlation is 0.71 and statistically significant.⁹ The high magnitude of this correlation suggests a positive association between the propensity of a firm to form alliances and the degree to which it innovates in technological fields that are not directly related to those in which it has developed technologies in the past (Stuart, 1995, presents more systematic evidence of this). The decision to branch out from a firm's existing fields of innovation is the most likely source of its movement in the configurations.

DISCUSSION: TECHNOLOGICAL POSITIONS, INNOVATIVE CAPABILITIES, AND STRATEGIC GROUPS

The core imagery that underlies this work is the conception of the technological base of an indus-

⁸ We coded patent license and cross-license, second source, joint ventures, joint product development, and technology exchange agreements for this analysis.

⁹ In a comprehensive data base on strategic technology alliances, Hagedoorn (1993) reported that the Mitsubishi Group had the highest total number of technology partnerships among all firms worldwide. Hagedoorn found that Mitsubishi formed 157 alliances in the 1980–84 period, and 293 alliances in the 1985–89 period. Hitachi and Toshiba were also among the world's 10 most frequent technology partners.




Figure 6. Technological positions of Japanese semiconductor firms and strategic alliances: 1982, 1987, and 1992

try as an evolving network. Discrete inventions, which for the most part belong to corporate innovators, form the nodes of this network. Technological commonalities among the inventions in the network are the ties that connect nodes. In an ongoing research program, we suggest that patents and patent citations can be used to represent this expanding 'technological network' for a select number of high-technology industries (Stuart, 1995; Podolny and Stuart, 1995). Collecting all patents in particular technological areas and aggregating firms' inventions allows the computation of network-theoretic attributes of firm-level positions in this technological network. Like the measures and methods used in this study, ideas and techniques in the network literature lend theoretical insights to the computation and behavioral implications of different attributes of firms' technological positions.

In the analysis of this paper, the proximity of two firms' positions depends upon the degree to which they are structural equivalents. Two actors who are perfect structural equivalents are also assumed to perform the same role in the relational

structure in which equivalence is measured (i.e., they are role equivalents). Generalizing this insight to the empirical context of this article, two firms that occupy structurally equivalent positions in the technological network do so because they perform similar roles as innovators. Two such firms would appear very near to one another in the positional maps of this paper. In principle, they could substitute for one another in their innovative roles.

Assume that the ability to develop the inventions that form the basis for a particular innovative role rests on the incumbent's accumulation of difficult-to-imitate innovative skills. This is not a Herculean assumption: following the discussion of the broad literature that characterizes innovation as a path-dependent process, it is quite plausible that firms' positions derive from skills that are in fact quite difficult for competitors to replicate quickly. Moreover, it is also the case that a well-honed innovative capability can be an extremely valuable resource to high-technology firms. In this case, the configurations of firms can be viewed as maps of innovative capabilities.

We believe that the analysis in this paper has an obvious link to the resource-based view of the firm: the configurations of Figures 1–6 represent one approach to positioning firms on the basis of inimitable, valuable resources that are potential sources of sustainable competitive advantage.

In addition, the analyses of this paper are pertinent to the literature on strategic groups. If positions cohere because the firms that hold them perform similar innovative roles, then clusters of firms can be viewed as grouping based on similar innovative capabilities. To date, scholars have usually identified intraindustry group structure by categorizing firms according to their product market positions, or else by general descriptors of their corporate strategy. However, as a number of scholars have suggested (McGee and Thomas, 1986; Dierickx and Cool, 1989; Barney, 1991), one point of contact between the resource-based view of the firm and work on strategic groups is to define intergroup mobility barriers in terms of heterogeneities among groups of firms in their possession of strategically valuable resources.

The findings of this paper do suggest that mobility barriers segregate technological positions. Furthermore, as other researchers have argued, the technological areas targeted by a firm's inventions in large measure circumscribe the expertise that it develops in manufacturing, marketing, and other core business functions (Teece, 1988). It is therefore compelling to use similarity of technological position as a basis for identifying groups whose capabilities are not easily imitated.¹⁰ To identify groups in the square matrices of interfirm technical proximity scores (the community matrices), one would apply a hierarchical clustering algorithm to partition the sample members.

¹⁰ It is important to note that intergroup mobility barriers may be asymmetric, even when firms' group affiliations are determined by their innovative capabilities. For example, hierarchical cluster analyses suggest that the firms in the region labeled 'CORE' in Figure 5 comprise one group, and those in the region labeled 'PERIPHERY' form a second group. The actual technological areas that comprise the basis of the inventive activities of the 'PERIPHERY' firms (e.g., linear ICs and discrete devices) are less complex than those that are the focus of the 'CORE' producers (e.g., optoelectronics and MPUs). In fact, most of the core firms produce linear ICs and discrete devices, in addition to more complex devices. Therefore, mobility barriers are asymmetric: it would be easier for the core firms to move into the periphery region than vice versa.

CONCLUSION: DIRECTIONS FOR FUTURE RESEARCH

The objective of this paper has been to develop a generalizable methodology for quantifying the evolution of firms' technological positions. Our approach conceptualizes the context of search in terms of the actual technologies developed by a sample of innovators, and the outcome of search in terms of its impact on firms' technological positions. From our perspective, an important and underemphasized component of the dynamics of technological change is that firms do not search in isolation; rather, they search as members of a population of simultaneously searching organizations. The methodology that we have suggested in this paper implicitly recognizes that a firm may come to occupy a differentiated technological niche not necessarily as the result of its own R&D, but as the result of the R&D of its competitors. In effect, a firm's position depends as much on the trajectories adopted by other firms as it does on its own trajectory.

A contribution of this research is that it offers a systematic conception of the context of search. The absence of such a method is surprising, particularly considering that the characterization of the search process is an essential step in the construction of evolutionary models of industry dynamics (Nelson and Winter, 1982; Winter, 1984). A direct consequence of the lack of a generalizable approach has been the inability to empirically test the basic assumption of local search in a convincing manner. For example, there have been no empirical tests of the Markovian assumption that the innovative direction of a company at period t_{m+1} depends critically on the state that it occupied at period t_m , but not on its prior history (Nelson and Winter, 1982). Additionally, it has been difficult to understand how the search environment and the history of search explain current technological positions and constrain future shifts in innovative directions.

An important influence on the findings of this study was the choice of setting—the semiconductor industry. Semiconductor technology is known to be cumulative, and because of its great complexity the technology is notably domain-specific. With few exceptions, the fact that a firm excels at innovating or producing in one market niche does not imply a similar expertise in a different niche. With the proposed methodology, however,

it would be possible to make intersample comparisons. Because the community matrices contain information on the global positions of all of the members of a system, metrics of the stability of the community matrices are comparable across systems of the same size. For example, it would be possible to compare the 10 largest Japanese firms to the 10 largest U.S. firms during the same interval of time to determine which group experienced the most change. It would also be possible to compare, for instance, the community matrices representing the 50 largest semiconductor firms to those representing the 50 largest pharmaceutical firms during the same time period. An analysis like this could assess the degree to which 'localness' characterizes the search trajectories of firms in different industries. A prior expectation would be that semiconductor firms are substantially more locally bounded in their innovation than pharmaceuticals because semiconductor technologies are more cumulative than are drug discovery techniques.

Innovation can be considered to encompass a broad array of technical and commercial functions, ranging from basic R&D to marketing. Our concern with firm trajectories in knowledge creation has led to our focus on the invention-generating stages of the innovation claim. However, we believe that the methodology that we have presented can be generalized to other stages of the innovation chain. For example, there exist several data sources that provide information on firms' participation in different product market niches in the semiconductor industry. Using such information and distance metrics like those employed in this paper, it is straightforward to measure the distance between firms in product space, just as we have measured the distance between firms in technology space. With this additional information, it would be possible to map evolving market positions and to explore the relationship between market and technological positions.

One of the most suggestive findings of the analysis is that Mitsubishi's movement into the 'technological core' was preceded by alliances with firms in that position. This association suggests the possibility that considerable shifts in technological position are facilitated by efforts to assimilate the technological developments of the firms in the areas to which a firm seeks to move. Alliances and acquisitions represent possible strat-

egies to bring about significant shifts in technological focus. One possible direction for future research would be to more systematically investigate the effect of alliance strategies and other strategic undertakings on the amount and direction of firms' search. In effect, one could investigate the impact of alliances or acquisitions on the distance of a firm's movement as specified by Equation 4. Similarly, it would be a simple extension of this research to model the effects of organizational characteristics—such as age or size—on the stability of a firm's technological position. The central question guiding this type of analysis would be: How do variables that proxy for the institutionalization of organizational routines affect the degree of inertia in the direction of firms' innovation?

REFERENCES

- Abernathy, W. J. and K. B. Clark (1985). 'Innovation: Mapping the winds of creative destruction', *Research Policy*, **14**, pp. 3–22.
- Albert, M. B., D. Avery, F. Narin and P. McAllister (1991). 'Direct validation of citation counts as indicators of industrially important patents', *Research Policy*, **20**, pp. 251–259.
- Barney, J. B. (1991). 'Firm resources and sustained competitive advantage', *Journal of Management*, **17**(1), pp. 99–120.
- Boeker, W. (1989). 'Strategic change: The effects of founding and history', *Academy of Management Journal*, **32**(3), pp. 489–515.
- Burgelman, R. A. (1994). 'Fading memories: A process theory of strategic business exit in dynamic environments', *Administrative Science Quarterly*, **39**, pp. 24–56.
- Burt, R. S. (1987). 'Social contagion and innovation: Cohesion versus structural equivalence', *American Journal of Sociology*, **92**, pp. 1287–1335.
- Burt, R. S. (1988). 'The stability of American markets', *American Journal of Sociology*, **88**(2), pp. 356–395.
- Cohen, W. M. and D. A. Levinthal (1989). 'Innovation and learning: The two faces of R&D', *Economic Journal*, **99**, pp. 569–596.
- Cyert, R. and J. G. March (1963). *A Behavioral Theory of the Firm*. Prentice-Hall, Englewood Cliffs, NJ.
- Dierickx, I. and K. Cool (1989). 'Asset stock accumulation and sustainability of competitive advantage', *Management Science*, **35**(12), pp. 1504–1514.
- Dosi, G. (1988). 'The nature of the innovative process'. In G. Dosi, C. Freeman, R. Nelson, G. Silverberg and L. Soete (eds.), *Technical Change and Economic Theory*. Pinter, London, pp. 221–238.
- Eisenhardt, K. M. and C. B. Schoonhoven (1996). 'Resource-based view of strategic alliance formation: Strategic and social effects in entrepreneurial firms', *Organizational Science*, forthcoming.

- Hagedoorn, J. (1993). 'Strategic technology alliances and modes of cooperation in high-technology industries'. In G. Grabher (ed.), *The Embedded Firm: On the Socioeconomics of Industrial Networks*. Routledge, London, pp. 116–137.
- Hannan, M. T. and J. H. Freeman (1989). *Organizational Ecology*. Harvard University Press, Cambridge, MA.
- Helfat, C. E. (1994). 'Firm specificity and corporate applied R&D', *Organization Science*, **5**(2), pp. 173–184.
- Henderson, R. M. and K. B. Clark (1990). 'Architectural innovation: The reconfiguration of existing product technology and the failure of established firms', *Administrative Science Quarterly*, **35**, pp. 9–31.
- Jaffe, A., M. Trajtenberg and R. Henderson (1993). 'Geographic localization of knowledge spillovers as evidenced by patent citations', *Quarterly Journal of Economics*, pp. 577–598.
- Kimura, Y. (1988). *The Japanese Semiconductor Industry: Structure, Competitive Strategies, and Performance*. JAI Press, Greenwich, CT.
- Kimura, Y. (1989). 'Firm-specific strategic advantages and foreign direct investment behavior of firms: The case of Japanese semiconductor firms', *Journal of International Business Studies*, pp. 296–314.
- Kruskal, J. B. (1964). 'Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis', *Psychometrika*, **29**, pp. 1–27.
- Langlois, R., T. A. Pugel, C. S. Haklisch, R. R. Nelson and W. G. Egelhoff (1988). *Micro-Electronics: An Industry in Transition*. Unwin Hyman, Boston, MA.
- Lee, D. M. S. and T. J. Allen (1982). 'Integrating new technical staff: Implications for acquiring new technology', *Management Science*, **28**, pp. 1405–1420.
- MacKenzie, D. (1992). 'Economic and sociological explanation of technical change'. In R. Coombs, P. Saviotti and V. Walsh (eds.), *Technological Change and Company Strategies: Economic and Sociological Perspectives*. Harcourt Brace Jovanovich, London, pp. 25–48.
- March, J. G. (1988). *Decisions and Organizations*. Blackwell, New York.
- McGee, J. and H. Thomas (1986). 'Strategic groups: Theory, research, and taxonomy', *Strategic Management Journal*, **7**(2), pp. 141–160.
- McPherson, J. M. (1983). 'An ecology of affiliation', *American Sociological Review*, **48**, pp. 519–532.
- Nelson, R. R. and S. G. Winter (1973). 'Toward an evolutionary theory of economic capabilities', *American Economic Review*, **63**, pp. 440–449.
- Nelson, R. R. and S. G. Winter (1982). *An Evolutionary Theory of Economic Change*. Belknap Press, Cambridge, MA.
- Podolny, J. M. and T. E. Stuart (1995). 'A role-based ecology of technological change', *American Journal of Sociology*, **100**(5), pp. 1224–1260.
- Podolny, J. M., T. E. Stuart and M. T. Hannan (1996). 'Networks, knowledge, and niches: Competition in the worldwide semiconductor industry, 1984–1991', *American Journal of Sociology*, forthcoming.
- Rosenberg, N. (1969). 'The direction of technological change: Inducement mechanisms and focusing devices', *Economic Development and Cultural Change*, **18**, pp. 1–24.
- Sahal, D. (1985). 'Technological guideposts and innovation avenues', *Research Policy*, **14**, pp. 61–82.
- Shan, W. (1990). 'An empirical analysis of organizational strategies by entrepreneurial high-technology firms', *Strategic Management Journal*, **11**(2), pp. 129–139.
- Stuart, T. E. (1995). 'Positions in a patent network: A theory of the rate and intensity of innovation', Working paper, Graduate School of Business, University of Chicago.
- Teece, D. J. (1988). 'Technological change and the nature of the firm'. In G. Dosi, C. Freeman, R. Nelson, G. Silverberg and L. Soete (eds.), *Technical Change and Economic Theory*, Pinter, London, pp. 256–281.
- Tushman, M. L. and P. Anderson (1986). 'Technological discontinuities and organizational environments', *Administrative Science Quarterly*, **31**, pp. 439–465.
- Tushman, M. L. and L. Rosenkopf (1992). 'Organizational determinants of technological change: Toward a sociology of technological evolution', *Research in Organizational Behavior*, **14**, pp. 311–347.
- Winter, S. G. (1984). 'Schumpeterian competition in alternative technological regimes', *Journal of Economic Behavior and Organization*, **5**(3–4), pp. 287–320.