Division of Research
School of Business Administration
The University of Michigan

October, 1991

A LANDSCAPE THEORY OF ALLIANCES WITH APPLICATION TO STANDARDS SETTING

Working Paper #666

Robert Axelrod*, Will Mitchell**
Robert E. Thomas**, Scott Bennett*
Erhard Bruderer**
The University of Michigan

*Political Science Department
**School of Business Administration

We are indebted to John E. Jackson and Roger C. Kormendi for their help in initiating this work. We appreciate comments given by Michael D. Cohen, John H. Holland, Valerie Y. Suslow, and Hal R. Varian, as well as advice received from participants at the Conference on Industrial Organization, Strategic Management, and International Competitiveness, held at the University of British Columbia in June 1991. We are grateful for information provided by Jim Linck, Mark Freed, Dave Martens, and Byron Askin of the University of Michigan MBA Program, and by several anonymous computer industry experts. This work was supported in part by National Science Foundation grants SES 8808459 and SES 9106371 to Robert Axelrod.

A LANDSCAPE THEORY OF ALLIANCES WITH APPLICATION TO STANDARDS SETTING

ABSTRACT

Landscape theory predicts how business firms (or other actors) will form alliances. The theory assumes that firms are myopic and their choices are incremental. The firms are myopic in the sense that they use information based only on pairwise relationships. The choices are incremental in that changes in alliance membership are made by only one firm at a time. Each partitioning of the firms into alliances results in a degree of frustration, and the predicted alliance configurations are those that have a local minimum of frustration. Landscape theory does well in predicting the choices of nine computer companies in joining one of two alliances promoting competing Unix operating system standards in 1988. The theory is applicable to many other aggregation processes as well as business alliances.

A LANDSCAPE THEORY OF ALLIANCES WITH APPLICATION TO STANDARDS SETTING

Alliances have played significant roles throughout economic history, often functioning as coordinating institutions in the creation and functioning of both barter and advanced industrial economies. Although many alliances arise among producers of complementary products and services, they also appear among competing firms. Examples of alliances among competitors occur throughout the economic sector. Investment banks form syndicates to market securities; insurance companies commonly form alliances to pool risk; and competing firms frequently join trade associations that lobby for favorable legislative and regulatory treatment. In the past decade, semiconductor and computer firms have formed R&D consortia such as Sematech to develop and manufacture new technology jointly. Product standards, such as the VHS standard for video cassette tapes, also are often set through the formal or informal allied actions of firms in an industry.

Although there are many types of economic alliances, economic research has focussed on alliances that coordinate market prices and outputs. A substantial literature has developed on the theoretical and empirical conditions for the existence and stability of cartels and other market coordination alliances.[1] There is also a smaller theoretical literature on the alliance composition question, that is, which firms will combine to form alliances.[2] Unfortunately, there has been no empirical test of any game theoretic alliance composition model.

The alliance composition question is especially important when competing alliances form. If one alliance becomes dominant, companies that choose a losing alliance will often suffer competitively. Moreover, because different firms have different capabilities, the identity of the members of an alliance will often influence the actions and output of the alliance. For example, the composition of an alliance may provide it with an advantage in getting a new product to market, but leave it with poorer technical design abilities than those held by a competing alliance. If the marketing capabilities dominate, a technologically inferior product may become a product standard. Therefore, not only will the types of products available to consumers differ depending on the composition of the alliances formed, but the dominant products may be technically inferior.

This paper augments the literature on alliance composition by introducing a landscape theory of alliance formation. The theory provides a new way of analyzing alliance composition problems and a method for predicting the composition of actual alliances. The landscape theory approach (Robert Axelrod, 1991) is sufficiently general to apply not only to economic coordination alliances but to many other aggregation processes as well, including international alliances, coalitions of political parties in parliaments, social networks, social cleavages in democracies, organizational structures and the structure of belief systems. We illustrate the application of landscape theory with the empirical case of membership in economic alliances formed to promote product compatibility standards for Unix computer operating systems.

I. Background

A. Composition of Alliances

Like landscape theory, the standard game theoretic alliance composition literature provides a theoretical framework for analyzing alliance composition problems. By calculating and comparing "coalition structure values" (Guillermo Owen, 1977) for each possible way of partitioning players into alliances (an alliance configuration), a standard game theoretic analysis can determine both the alliance configuration that is likely to emerge and the stability of each configuration. The configuration may be a single alliance or consist of competing alliances.

Unfortunately, the coalition structure value framework is difficult to apply to empirical data because of its substantial information requirements. Empirical application of this literature would require identifying and quantifying payoffs for each participant in every conceivable set of alliances, which is a daunting task for any analysis. Using this literature to analyze complex alliance composition problems is especially difficult because payoffs for each firm depend upon the choices made by all other firms, rather than upon just pairwise relationships. Consider the example of the standards setting case. The size of the market will vary with the number of standards, while a given firm's market share will vary with (among other factors) how quickly the firm can bring a product to market relative to other firms. How quickly each firm develops a product depends on several factors. These include how closely other alliance members cooperate with the firm, whether one or more members has a technological advantage or head start in producing related products, and whether one or more members is powerful enough to influence the selection in favor of a proprietary standard (Michael Katz and Carl Shapiro, 1986; H. Landis Gabel, 1987).

Thus, for complex alliance composition problems, it is virtually impossible to determine complete payoff functions as game theory traditionally requires. Landscape theory exploits the information present in simple firm attributes to avoid the use of payoffs and the difficulties of their measurement. Examples of attributes for a given firm include its pre-alliance market share, product specializations and expertise, and geographical or niche product market activity. Because it is possible to identify firm attributes and then concentrate on the various ways that the players could form alliances based upon their pairwise relationships, landscape theory can make empirical predictions based on much less information than the standard game theoretic approach requires.

B. The Landscape Approach

In an economic context, landscape theory starts with the premise that a firm joins an alliance in order to improve its performance. Specifically, firms will work well together if they perceive that allying will improve their market position, either by creating market power or by aligning complementary capabilities.[3]

Landscape theory makes two basic assumptions, both drawn from the recognition that it is difficult for a firm to assess the value of each potential alliance. The first assumption is that a firm is myopic in its assessments, by which we mean that a firm evaluates how well it gets along with any other firm independent of all the other members of an alliance. Therefore, a firm avoids the difficult problem of assessing all alliance members at once by making pairwise evaluations. The evaluations depend on firm attributes. For example, an alliance between a large retailer and a large manufacturer may benefit both firms, while two manufacturers are more likely to conflict in an alliance. Differences and similarities in attributes determine how strongly a firm wants to work with another firm.

Under landscape theory, a firm decides what alliance to join by assessing how well its own attributes complement or conflict with the attributes of each other firm in each alliance. There is often no way companies can make alliances such that all firms ally only with firms they want to be allied with and are separated from firms they want to be separated from. Typically, a firm will end up allied with some firms with which it would rather not be allied and separated from others with which it would prefer to be allied. When this occurs, firms are frustrated. Firms want to minimize the amount of frustration they incur. The landscape theory behavioral assumption that corresponds to payoff maximization is frustration minimization.

The second basic assumption of landscape theory is that adjustments to alliances will take place by incremental movement of individual firms. This rules out the possibility that a coalition will form within an alliance and then switch allegiance as a block. This strong assumption is appropriate when information regarding payoffs is uncertain, because the causal ambiguity between alliance actions and payoffs will both raise negotiation costs and reduce the ability of the firms to use side payments to arrive at an optimal solution.

Landscape theory uses several abstract concepts from the physical sciences and biology that have proven useful in studying the many possible ways in which elements of a system can fit together in more or less frustrated arrangements.[4] Using the theory, we calculate the energy of an alliance configuration, where an alliance configuration is a specification of which firms are in which alliance. A measure of the energy of an alliance configuration specifies the degree to which a given outcome leaves the firms frustrated, either because they are allied with firms that they would rather not be allied with or because they are separated from firms with which they would prefer to be allied. It is not possible to eliminate all frustration in most cases, but it is possible to minimize it. Landscape theory predicts which alliance configurations may occur, how much frustration is inevitable, and how the system will respond to changes in the relationships between the elements.

The basic idea of landscape theory is that each configuration of firms partitioned into alliances can be represented as a point in an abstract space. Similar configurations are represented in this space as nearby points. Associated with each configuration is a number representing just how well that configuration fits together. This allows the energies of the configurations to be visualized as a landscape. An incremental change that lowers frustration can then be viewed as a downward movement on this surface.

C. Standards Setting

Standards setting provides a useful illustration and test of landscape theory. Like market coordination agreements, cooperating firms generally obtain significant benefits when compatibility standards are established. These gains exist because many firms and consumers are willing to adopt a technological product that has many other market participants committed to it (Katz and Shapiro, 1985, 1986).[5]

Establishing a compatibility standard often requires coordinated action between competing companies, particularly "when there is no clear [market] leader and there are different preferences among standards" (Joseph Farrell and Garth Saloner, 1988 p. 236). However, setting standards may be more difficult than coordinating markets because of two conflicting objectives. First, every firm wants its own proprietary technology to prevail to give it a competitive edge. Second, every firm wants as few standards as possible to survive, in order to maximize the size of the market. In the market coordination problem, an agreement often can be renegotiated to address a failing with the initial agreement. By contrast, once an industry agrees on a standard, it is exceedingly difficult to alter that standard (Farrell and Saloner, 1985). Thus, alliance formation is especially critical to standards setting because once the alliance chooses a standard, there may not be a later opportunity for a rejected technology to become a standard. The standards setting case also differs from the market coordination problem in that the act of setting standards is at least one step removed from allocating market shares. Therefore, unlike the market coordination problem, the standards setting problem has no unambiguous causal link from alliance choice to market payoff.[6]

Our empirical study is of the struggle by technical workstation firms to create a Unix operating system standard. The technical workstation market, which emerged during the 1980s, has depended on several incompatible implementations of the Unix operating system. The attempt by two companies to impose a Unix standard failed utterly. Out of this failure, two opposing alliances formed in 1988. We use two simple attributes of firms and their 1987 market-related sizes to predict commitment of each relevant firm to the standards promoted by the two alliances. As we will show, this application of landscape theory correctly predicts the early commitment to the competing standards made by all but one firm in our sample. The study also provides insight into major factors that influenced various player's choices.

II. Landscape Theory

The primary concepts of landscape theory are propensity, size, and distance. The key assumption of landscape theory is that each pair of firms, i and j, has a propensity, p_{ij}, to work together. Propensity is a measure of how willing the two firms are to be in the same alliance together. The propensity of two firms to ally is positive if the two firms will get along well together. The propensity is negative if they will have difficulty working together. For example, suppose that the more dissimilar two firms are the easier it would be for them to work together because their firm-specific assets complement each other, while the more similar they are the harder it would be for them to work together because their firm-specific assets are inherently competitive. In the simplest case where each complementary attribute has equal importance, the propensity of any two firms to ally could be measured as the number of relevant differences they had minus the number of similarities. For present purposes it is convenient and plausible to assume that propensity is symmetric, so that $p_{ij} = p_{ji}$.

The size of a firm, s_i, reflects the importance of that firm relative to the others. Size might be measured by market share, investment resources, or other considerations, depending on what is important in a particular application.

An alliance configuration specifies which firms are in which alliance, and so determines the distance between any two firms, i and j. In the simplest version of the theory all firms are in one of two possible alliances, so we can let distance be 0 if the two firms are in the same alliance, and 1 if the two firms are in different alliances. In different situations, more than two alliances could be allowed and other measures of distance in a configuration could be used (Axelrod, 1991). In an organization, for example, the distance between any two people could be taken to be the sum of the distances to their lowest common superior.

Using propensity, size, and distance, it is possible to define a measure of how poorly or well a given configuration satisfies the various propensities of firms to be near or far from each other. This measure is called the energy of an alliance configuration. Energy (or total frustration) is low if the propensities are well satisfied and high if they are not. The formula defining energy of a configuration is

[1] $$E(X) = \sum_{i,j} s_i s_j p_{ij} d_{ij}(X)$$

where s_i and s_j are the sizes of two firms, p_{ij} is the propensity of the two firms to be allied with each other, and $d_{ij}(X)$ is distance between the two firms in configuration X. The intuitive idea behind this formula is that energy is lower (and the configuration is less frustrating) when firms that want to work together are in the same alliance, and those that want to compete are in different alliances. Size plays a role because having a proper relationship with a large firm is more important than having a proper relationship with a small firm.

Once energy is defined, the abstraction begins to pay off. Given the energy of each configuration, it is possible to construct an energy landscape. The landscape is simply a graph with a point for each possible alliance configuration, and a height above this point for the energy of that configuration. Figure 1 shows an example of a landscape where each point at the bottom of the figure indicates a specific configuration, and the vertical dimension represents the energy of that configuration.

**** Figure 1 here: A Landscape With Two Local Optima ****

The landscape allows predictions to be made about the dynamics of the system. The basic idea is wherever the system starts it will tend to run "downhill" to adjacent configurations that have lower energy. Adjacent configurations are those which differ by the alliance membership of a single firm. Given that movement will be downhill to adjacent configurations, the changes in alliance patterns will stop only when the partners reach a configuration that is a local minimum. For example, there are two local minima in Figure 1, and therefore two potentially stable configurations. Stability here means that the configuration (i.e., alliance pattern) does not change any further. Since adjacent configurations differ by a single actor, the locally optimal configurations are Nash equilibria (i.e., points from which no individual would move unilaterally).

Which stable configuration will occur depends on where the system starts. For example, most configurations on the left half of Figure 1 will move to the low point on the left side, which happens to be the global minimum. The set of configurations that would lead to a given minimum can be thought of as the basin of attraction of that minimum.

A characteristic of landscape theory is that there may not be any configuration that completely satisfies everyone. Even for a configuration that is the global optimum, most or even all firms may be somewhat frustrated in the sense that some "friends" will be in the opposite alliance or some "enemies" will be in the same alliance.

A further implication is that the local optimum into which the system settles can depend on the history of the system. Early history, which might be in part the consequence of small events or even chance circumstances, can determine which outcome prevails. Thus, for better or worse, the outcome may be in part the consequence of a "frozen accident" just as the QWERTY keyboard is (Paul David, 1985; S.J. Liebowitz and Stephen Margolis, 1990). Once the system settles into a basin and moves toward a local minimum, it may be difficult to leave that basin. Moreover, if there is more than one local optimum, the one that the system moves toward may not be the one with the least total frustration.

III. The Case Of Unix Standards Setting Alliances

A. Technical Workstations And Unix Alliances: Historical Background

The struggle over Unix standards for technical workstations provides an illustration and test of landscape theory.[7] Technical workstations are powerful desktop computers, typically used in engineering and scientific applications. The first commercial technical workstation was introduced about 1980. Worldwide sales were $2.5 billion in 1987 and reached $10 billion in 1990. A key aspect of technical workstation design is the operating system, which controls the hardware and manages the flow of information and communication among the various components of the computer system. Most commercial technical workstations have used some version of the Unix operating system, which the American Telephone and Telegraph Company (AT&T) developed at its Bell Laboratories during the 1960s. Altogether, more than 250 versions of Unix-based operating systems have been designed by computer hardware companies and academic institutions. Applications software written for one version often does not operate on another Unix system.

In 1984, the software incompatibility across Unix operating systems induced several leading European, American, and Japanese computer manufacturers to form the X/Open group with the goal of encouraging the development of Unix standards. This consensus approach to standardization failed in October 1987, however, when X/Open members AT&T and Sun Microsystems, Inc. announced that they would pursue development of a Unix operating system based on AT&T's System V. The new system would be available to other companies under proprietary license.

AT&T and Sun did not proceed unchallenged. Seven major computer manufacturers, including the Digital Equipment Corporation (DEC) and International Business Machines Corporation (IBM), responded in May 1988 by forming the Open Software Foundation (OSF). A primary purpose of the OSF was to develop a standardized Unix operating system that did not draw on AT&T's proprietary technology.

AT&T and Sun responded to the OSF by forming Unix International, Inc. (UII) in December 1988, the members of which would advise and promote the development of AT&T's Unix System V. Although some secondary participants joined both alliances, there was no overlap among the nine full sponsors of the OSF and the 10 principal members of UII.

Both UII and the OSF have expanded since 1988 and each has continued its efforts to create an operating system standard. AT&T introduced a commercial release of System V version 4 in November 1989 and the OSF released a commercial version of OSF/1 in late 1990. These systems compete for acceptance by computer buyers and standardization committees throughout the world.

This chronology follows the pattern of Joseph Farrell and Garth Saloner's (1988) argument that adoption of compatibility standards may be achieved by means of a hybrid of committee coordination and market leadership. We can view X/Open as an attempt at committee coordination, the AT&T-Sun action in 1987 as a unilateral market leadership move, and the formation of the OSF and UII as subsequent attempts to create standards by committees. The AT&T-Sun attempt at market leadership did not succeed because the two firms were not strong enough to exert their will on other strong firms such as DEC and Apollo, and because there was no consensus that AT&T's new Unix operating system would succeed. Therefore, computer manufacturers created two competing alliances to promote Unix standards.

B. Landscape Theory in the Unix Case

The first task in applying landscape theory is to identify the relevant firms. We chose the companies that had the greatest influence on the technical workstation market in 1987, the year before the OSF and UII alliances formed.[8] We included the eight firms having at least one percent of the worldwide technical workstation market in 1987, together accounting for over 95% of market revenue. We also included AT&T because it held the copyright for the parent version of Unix, licensed Unix to technical workstation manufacturers, continued to develop Unix in 1987, and was a potential entrant to the technical workstation industry.

To calculate the energy of potential alliances among the nine firms, we required measures for size and propensity. For the size of the eight technical workstation manufacturers, we used each firm's 1987 share in the technical workstation market. We judge this to be an appropriate measure of the firm's importance in a standard-setting alliance.[9] To assign a size for AT&T, which did not manufacture technical workstations in 1987, we asked four computer industry experts to estimate the importance of AT&T with respect to its influence in establishing a Unix standard. We assigned a size weight of 28.5 based on the median of the experts' estimates. The size places AT&T among the most important firms in establishing a Unix standard.[10]

To choose measures of the firms' propensities to collaborate in the Unix standards setting alliances, we first decided that the most important general issue was the intensity of competition between the companies. We then identified factors associated with the intensity of competition in the technical workstation industry. We found two important competitive factors that characterized the technical workstation industry in 1987. First, the ability to design innovative products was an attribute that distinguished firms in the industry. Firms with design capabilities competed with other design firms on the basis of innovation, whereas nondesign firms (assemblers) either copied or purchased the innovations created by design firms. Therefore, we expect two designers to compete and two assemblers to compete, but a designer and assembler to have a reason to cooperate.

Second, some computer companies specialize in the workstation market and others are generalists with many lines of computer-related products. We expect two generalists to compete with each other in many markets and two technical workstation specialists to compete with each other in the workstation market. Because of the head-to-head competition in several markets by generalists and in the technical workstation market by specialists, we expect pairs of similar firms to find it difficult to work together in a standards setting alliance. However, we expect a generalist and a specialist to have reasons to cooperate in technical workstation commercialization, with the generalist providing distribution capabilities and financing while the specialist provides technical skills.[11] Of the nine firms in this study, five were workstation specialists and four were generalists. All nine possessed design capability.

We used the competitive attributes to calculate propensities for each firm to ally with each other firm.[12] We assigned a value of 1 when two firms had complementary attributes (specialist-generalist) and -1 when attributes conflicted (two specialists, two generalists, or two designers). We then summed these values to create propensity measures for each pair of firms. Thus, the propensity of two firms to work together was either 0 (when one firm was a specialist and the other a generalist, and both were designers) or -2 (when both were specialists or generalists, and both were designers).

Together with size, we expect the propensities to predict the commitment of firms to developing one of two potential standards. The extent to which the coalitions predicted by landscape theory conform to the actual commitment to the operating systems promoted by the OSF and UII is a test of the power of the theory. For the Unix case, we assume that there will be at most two alliances, as has occurred, rather than trying to predict the number of alliances that will form. This assumption is appropriate for many standard-setting cases, because the benefits of standardization will decline as the number of designs increases.[13]

C. Predictions

The predictions generated by the landscape analysis are displayed in Table 1, along with the attributes and size of each firm. With nine firms, there are 2^8=256 possible alliance configurations. Two configurations of alliances had locally minimum energy. We refer to these as configurations 1 and 2. These two configurations had equal energy and so were tied for global optimum.[14] The two configurations are quite different. Sun and Prime join with IBM and AT&T in a locally optimal alliance in configuration 1, but are separated from them in configuration 2. DEC and HP ally with Apollo, Intergraph, and SGI in configuration 1, but are separated from them in configuration 2.

**** Table 1 here -- The Two Configurations Predicted By Landscape Theory Analysis ****

To see how the energy of the alliance configurations is calculated, consider configuration 1 in Table 1. In this configuration Sun is in an alliance with Prime, IBM, and AT&T, while the remaining five companies are in a competing alliance. The total energy of configuration 1 is calculated by summing the energy contributed by each distinct pair of firms.

Let us start with Sun's pairing with each other firm.

1. The pairings with the other three firms in the same alliance as Sun (whether generalist or specialist) contribute no energy because their distance from Sun is zero.

2. The pairings with the two generalists in the alliance that excludes Sun (i.e., DEC and HP) contribute no energy because their net propensities with Sun are zero (the positive generalist-specialist attribute cancels the negative design-design attribute).

3. The pairings with the three specialists in the alliance that excludes Sun (i.e., Apollo, Intergraph, and SGI) have two negative attributes, because Sun shares both the specialist and design attributes with each firm. Thus, the propensity of Sun to work with each of them is -2. Because Sun is not in the same alliance, Sun is at unit distance from them. To illustrate the calculation that equation [1] requires, consider the contribution to system frustration of Sun and Apollo when they are in the same alliance. The contribution of the symmetric Sun-Apollo and Apollo-Sun pairs is two times the size of Sun times the size of Apollo times the propensity of Sun and Apollo to work together

times their distance from each other in configuration 1. This product is (2)(28.9)(21.1)(-2)(1) = -2450.72. Similarly, the total pairwise contribution of Sun and Intergraph is (2)(28.9)(4.4)(-2)(1) = -508.64. Because SGI and Intergraph are the same size, the total pairwise contribution of Sun and SGI is also (2)(28.9)(4.4)(-2)(1) = -508.64.

To calculate the total energy of configuration 1, the same type of reasoning is used for each other pair of firms. The result for configuration 1 is -7657.80. This configuration is a local optimum because all the configurations that differ by a single firm have greater energy.

There is only one other local optimum, configuration 2. As we noted above, the two locally optimal configurations have equal energy and therefore tie for global optimum. Our knowledge that the first move in the alliance forming process was AT&T joining with Sun allows us to predict that the result will be configuration 1 rather than 2. We are predicting whether each of the other seven firms will commit to the same or different operating system as AT&T and Sun, based on configuration 1.

D. Results

We judged a firm's commitment to development of an operating system by identifying the first of the OSF and UII alliances it joined. Apollo, DEC, IBM, and HP were founding members of the OSF in May 1988, while Intergraph and SGI joined in late 1988. ATT, Sun, and Prime were founding members of UII in December 1988. We obtained information about alliance membership from the firms' 10-K reports to the Security and Exchange Commission, from articles in the business press, and through conversations with individuals at OSF, UII, and several of the companies.

Table 1 shows the predicted commitments in configuration 1 compared to what happened. In configuration 1, six of the seven commitments made after the AT&T-Sun link-up in 1987 are predicted correctly. Only IBM is incorrectly predicted. The close match between empirical commitments and the predicted alliances provides support for landscape theory. If we start with the knowledge that AT&T and Sun will be allied, then there are $2^7=128$ ways of assigning the remaining seven firms to the alliances. Of the 128 possible configurations, one is completely correct and seven are off by one firm. Therefore, the probability of getting seven or more of the eight right by chance is (1+7)/128=1/16.

Because the industry experts' estimates of AT&T's size varied substantially, we tested the sensitivity of the results to the estimate.[15] Knowing that AT&T joins Sun, the predictions are off by one firm if the AT&T size is set from 9 to 31, inclusive. There are two errors with AT&T set at 8 or less. With an AT&T size of 32 or more, there are no misassigned firms.

The changes in the predictions that occur as AT&T's size varies reflect plausible strategic considerations. Because AT&T is a generalist, the predicted division of the five specialist firms does not depend on AT&T's size and is always correct. However, the predicted division of the four generalists is sensitive to AT&T's size. Two generalists, AT&T and DEC, always oppose each other, but HP and IBM shift sides depending on AT&T's importance. If AT&T is small (1 to 8), then HP and IBM ally with AT&T to counter DEC. HP continues to ally with AT&T if AT&T has moderate size (9 to 16), but IBM moves to a partnership with DEC. If AT&T is larger (17-23), then either HP or IBM allies with DEC while the other allies with AT&T. If AT&T is large (24 to 31), HP allies with DEC while IBM allies with AT&T. Finally, if AT&T is very large (32 or more), both HP and IBM ally with DEC. The empirical outcome is consistent with the last case, that is, that AT&T was expected to play a very large role in setting the Unix standard.[16]

Another way of looking at the relationship between the theory and the size estimates is to ask what the theory would say about the size of AT&T given the alliance configuration that actually occurred. The answer is that the theory would suggest that

AT&T was very important (32 or more) because that size would accurately predict the alliance choices of each of the firms after the initial Sun-AT&T linkup.

IV. Conclusions and Future Research

The landscape theory approach to predicting alliance compositions has five key features, most of which are illustrated by the empirical application to Unix standards reported in this paper.[17] First, several distinct local optima may occur in an alliance composition problem, because some landscapes will contain more than one point at which any incremental change would increase the frustration of the system (W. Brian Arthur, 1988, 1990). In the Unix case, we found two different stable alliance configurations. Four firms would have had to change alliances to transform one of these stable configurations into the other. Second, some stable configurations may be inefficient, in the sense that they are not the global optimum. However, both outcomes reported in this paper are efficient because they attained equal energy and so were both global minima.

The third key element of landscape theory is that firms become locked into a stable alliance, even if they are joined with strong competitors and opposed to firms with complementary capabilities. DEC and HP, which compete in several computer markets, were correctly predicted to be members of the same Unix standards alliance, opposing firms against which they compete in fewer markets. Fourth, outcomes are path-dependent because the specific configuration at which the system finds stability will be strongly influenced by the early moves in the alliance-building process.[18] For example, configuration 1 in Table 1 was closest to the observed outcome because AT&T and Sun allied together. Had another pair come together early in the process, say AT&T and Apollo, then a quite different outcome might well have emerged, such as configuration 2. The fifth feature is that it is difficult to predict which stable alliance configuration will emerge, because it is difficult to forecast which initial move will occur.

Although which of the two stable outcomes would emerge may not have been predictable, the features shared by the two potential outcomes helps us understand the empirical process of alliance formation in standards setting cases. As can be seen in Table 1, the combination of specialist-generalist attribute and total subgroup size is the key factor influencing the formation of the potential alliances. Both configurations divide the specialists along the same lines into two subgroups of almost equal size (Sun and Prime versus the other three specialists). Similarly, both configurations divide the generalists along the same lines into two subgroups of almost equal size (DEC and HP versus the other two generalists). Thus, firms may balance the conflict between enjoying the benefits of standardization and incurring the problems of associating with close competitors by splitting into groups that are as close as possible to equal size. Moreover, the fact that the specialists and generalists could both be divided into nearly equal-sized subgroups -- which will be common in industries with several moderately strong competitors -- may help explain why the Unix standards battle has been a standoff until now.

The analysis also illustrates how both history and expectations may influence the determination of a stable outcome when there are multiple possible equilibria (Katz and Shapiro, 1985; Krugman, 1991). History can play a strong role because firms may use past and current strengths to assess each other. The propensities and most sizes used in this study are based on historical positions in the technical workstation industry. Moreover, the initial move by AT&T and Sun determined which of the two stable outcomes would emerge. However, expectations may also play an important role. In this study, AT&T's size is based on expectations of its potential importance as a standard setter, even though it was absent from the 1987 technical workstation market.

Landscape theory complements the intellectual underpinnings of several current approaches to understanding responses to new economic conditions, such as technical change. The theory parallels studies of complex adaptive systems, which attempt to predict the choices that economic agents will make when confronted by a "novel and evolving world" (John Holland and John Miller, 1991 p. 365). With its incremental search for optimal alliance configurations, the landscape approach also complements Richard Nelson and Sydney Winter's (1982) evolutionary economic theory, which predicts that individual firms will incrementally adopt new routines that are similar to their established routines. Landscape theory provides a myopic complement to established game theoretic approaches to predicting which firms in an industry are likely to ally, because it requires only the data likely to be used by managers faced with limited information.

Finally, landscape theory can be used as a guide to corporate strategy and public policy. It can suggest where leverage can most efficiently be applied to move from one potential alliance configuration to another, or which initial move is likely to lead to a desired alliance. The theory could also identify which strategic coalitions are inherently unstable, and which would become stable if critical subgroups could be induced to change their allegiances.

In this paper we have described the current scope of landscape theory and a specific application of it, but the theory requires further development to guide consistent application to other cases. A key task is to develop a general method for determining what individual attributes are appropriate to use for propensity measures in any given case, or possibly determine direct measures of pairwise propensities that could be used instead of basing the comparison on individual attributes. In most commercial cases, we expect that firms possessing complementary capabilities will tend to get along well together, while those possessing similar capabilities will not. However, these relationships are not universal. If similar firms can create sustainable market power, as may occur in industries with a few strong players, they will have incentives to form a cartel. The common denominator is that attraction or conflict will be determined by the balance of the advantages of working together versus tendencies to compete. If the joint performance of two firms can be increased by a partnership, whether due to vertical integration or horizontal power, the firms will tend to get along well together. Although determining propensities will require substantial understanding of specific empirical circumstances, the primary commercial attributes are likely to involve historical, technical, and competitive factors. Competitive factors include both product and geographic competition.

Much important work also remains to be done to increase the scope of the theory. The following five issues deserve particular attention. (1) Being able to predict the number of alliances, as well as the composition of a given set of alliances, would increase the breadth of the theory.[19] (2) Allowing a firm to choose either neutrality or dual membership in alliances, as several firms have done in the Unix standards case would extend the theory.[20] (3) Allowing asymmetric propensities would extend the theory to instances in which only one firm in a pair is attracted to the other. (4) Allowing the sizes and propensities to change endogenously, and (5) allowing occasional nonincremental movement from one basin of attraction to another also would broaden the theory. Even now, though, this paper demonstrates the accuracy and power of landscape theory.

To understand the contribution of landscape theory, an analogy to research on the Prisoner's Dilemma is apt. The value of the Prisoner's Dilemma is not only that it gives accurate predictions, but also that it leads to a deeper understanding of strategic processes such as the way in which the shadow of the future is essential for cooperation among two competitors. Likewise, the value of landscape theory is not only in providing accurate predictions, but also in leading to a deeper understanding of alliance formation processes among many competing and complementary firms. Just as the Prisoner's Dilemma helps us see important similarities across a wide range of applications, landscape theory may help us see how alliance formation processes in many different applications are similar when viewed with the help of a common theoretical framework.

FOOTNOTES

1 Research on market coordination dominates the economic literature related to alliances. Early research on market coordination agreements concentrated on the alliance maintenance issue (e.g., Edward Chamberlin, 1933; George Stigler, 1964). This interest in the sustainability of market coordination agreements has continued with the application of repeated games (James Friedman, 1977; Ariel Rubinstein, 1979; Drew Fudenberg and Eric Maskin, 1986) and uncertainty theory (Edward Green and Robert Porter, 1984). Work by Arthur Fraas and Douglas Greer (1977) exemplifies empirical attention to market coordination issues.

2 The game theoretic alliance composition literature addresses the questions of whether an alliance will form (Reinhard Selton, 1973; Gregory Werden and Michael Baumann, 1986), and what will be the composition of alliances formed among a given set of players (Lloyd Shapley and Martin Shubik, 1969; Guillermo Owen, 1977; Sergiu Hart and Mordecai Kurz, 1983; Roby Rajan, 1989).

3 In other applications, the definition of "working well together" will correspond to the application. For example, a political coalition works well together if it increases the probability of electing a particular candidate or passing desired legislation.

4 The idea of an abstract landscape was originally developed to study potential energy in physical systems and had its first rigorous development in the context of Hamiltonian systems (Vladimir Arnol'd, 1978; Ralph Abraham and Christopher Shaw, 1983; Gregoire Nicolis and Ilya Prigogine, 1989). Biologists have independently developed landscapes to characterize evolutionary movement in an abstract "fitness landscape" of genes (e.g., Sewell Wright, 1932, Stuart Kauffman, 1989). More recently, energy landscapes have been used in artificial intelligence to characterize the dynamics of complex systems such as neural networks (John Hopfield and David Tank, 1986). (In biology and artificial intelligence, the polarity of the landscape is reversed so that the

improvement is thought of as hill-climbing rather than descent into valleys.) What is new in this paper is not energy or landscape, therefore, but using pairwise relationships based on attributes of the firms to predict alliance composition.

5 Stanley Besen and Garth Saloner (1989), Gary Hamel, Yves Doz, and C. K. Prahalad (1989), and Thomas Jorde and David Teece (1990) assert that competitors can often achieve substantial gains by cooperating in the development of new compatible technology. See Paul David and Shane Greenstein (1991) for an annotated bibliography of the role of compatibility standards.

6 In a popular industrial organization economics text, F.M. Scherer and David Ross (1990 p. 608) conclude that "little is known about the distribution of benefits and costs from ... standard-setting activities under diverse market structural conditions."

7 The historical information in this subsection comes from public sources, e.g., Computer Technology Research Corp. (1990a,b). Garth Saloner (1990) describes these events in more detail.

8 We focused on the technical workstation market because there was low cross-elasticity of demand between it and other computer markets in 1987, partly because the lack of standardized operating system made technical workstations difficult to sell in the nontechnical markets. One of the factors driving the attempted standardization was the desire of the technical workstation manufacturers to expand into nontechnical markets.

9 The technical workstation market shares used in this analysis are based on figures reported by the International Data Corporation, Dataquest Inc., and the business press (e.g., Kristiina Sorensen, 1988; Steven Burke, 1988).

10 The four experts were familiar with the technical workstation industry in 1987. One of the experts is a software firm executive. The second is an analyst with a computer software firm. The third is an industry analyst who specializes in Unix applications. The fourth expert is also a computer industry analyst. We sent each expert a letter in which we stated that we were interested in AT&T because it owned the copyright for the parent

version of Unix, licensed Unix to technical workstation manufacturers, continued to develop Unix in 1987, manufactured computer hardware, and was a potential entrant to the technical workstation industry. We then asked the experts in telephone conversations to estimate the importance of AT&T, relative to the market shares of the other eight firms in our sample.

11 For example, Sun agreed to share its expertise in designing workstations and microprocessors with AT&T partly in exchange for an infusion of over $300 million.

12 In addition to the competitive attributes, a technical attribute may influence the propensities to form alliances in the technical workstation market. Each technical workstation is designed around a particular central processing unit, including several types of complicated instruction set (CISC) and reduced instruction set (RISC) processors. One might expect firms that use the same processors to be likely to ally. The impact of the technical processor attribute was weak in 1988, however, because technical workstation design was significantly unsettled by the recent successful introduction of RISC-based systems. Most firms were at least considering incorporating new processors and so were less bound by past decisions than one might at first expect.

13 By allowing one or two alliances to form we differ from much of the traditional alliance literature, which assumes that there will be a single alliance or none. In practice, there are many examples of competition between two principle designs, such as the DOS and CP/M operating systems for personal computers, and VHS and Beta formats for video cassette tapes. The theoretical issue here is that the number of competing standards will be limited by the presence of network externalities (Katz and Shapiro, 1985), that is, cases in which the utility that a user derives from consumption of a good increases with the number of other consumers of the good. The tendency for positive consumption externalities to lead to a single standards setting alliance will often be countered, at least initially, by the desires of competing firms to influence and benefit from the standards setting process.

14 The dynamics along the surface of the predicted landscape are not directly comparable to the empirical process of alliance formation. The landscape theory approach calculates the energy of all possible alliance configurations in which each firm is committed to one alliance. In practice, by contrast, there was an initial pool of unallied firms that chose to commit to development of one of the two Unix standards in 1988.

15 The experts assigned AT&T scores of 8, 18, 39, and 50. Based on a presurvey decision, we used the median, 28.5. The two experts who assigned the lower estimates based their estimates primarily on AT&T's historically weak presence in the technical workstation market. The experts who assigned the higher estimates based their judgements on the company's 1987 potential to play a strong future role in creating a Unix standard.

16 We also generated predictions using propensities based on single attributes. In both cases -- the product-type and design attributes -- many stable configurations occurred (as many as 13), because neither attribute alone contained enough information to produce a more restricted set of predictions.

17 This list of landscape theory features is adapted from a discussion with Brian Arthur.

18 Only starting points on the boundary separating two or more basins will lead to more than one possible stable configuration. W. Brian Arthur, Yu. M. Ermoliev, and Yu. M. Kaniovski (1987) discuss the issue of path-dependence, that is, of the role of history in influencing outcomes. Path dependence also arises in the literature on trade with external economies (e.g., Paul Krugman, 1981; Wilfred Ethier, 1982; Arvind Panagaryia, 1986).

19 The following approach might be used to predict the number of alliances. First, let the size of the total market be a function of the number of competing standards and the market success of each standard. Then, measure the value to a firm of being in a alliance in terms of (1) the proportion of the total market an alliance would be likely to

capture in the future, and (2) the proportion of the alliance members' sales held by the firm. Market size would decline with increases in the number of alliances and increase with rising inequality among the market shares held by the alliances, holding other factors constant. The maximum market size would occur with a single alliance, but the incentive to form a single alliance would be balanced by the potential that a firm might dominate a smaller alliance. (Farrell and Saloner (1986) show that a firm with a relatively large installed base can dominate a rival with a smaller installed base.) If strong firms made two assumptions -- first, that the standard promoted by their alliance would emerge as the sole or dominant market standard; second, that they could capture a large proportion of the new sales generated when a dominant standard emerged -- they would sometimes find it worthwhile to maintain a competing standard-setting alliance until a dominant standard emerged.

20 At least two of the smaller firms in our sample eventually developed significant relationships with both UII and the OSF, and many of the other firms that have joined the alliances hold membership in both groups. In most cases, the dual memberships are held by relatively minor participants in the computer market, which will have little influence on the standard setting process. Although this point is outside the theory, it is not surprising to find that weaker players attempt to position themselves to adapt to whichever standard emerges. The small market share players have little influence on the configurations predicted by landscape theory. When we drop the smallest player (Prime) from the sample, for instance, we obtain the same predicted alliance configurations for the other firms as the configurations predicted when using all nine firms. This issue would also arise if the theory were applied to market coordination alliances, where many smaller firms are often not included in an organizing cartel.

REFERENCES

Abraham, Ralph H. and Shaw, Christopher D. *Dynamics: The Geometry of Behavior.* Santa Cruz, CA: Ariel Press, 1983.

Arnol'd, Vladimir Igorevich. *Mathematical Methods of Classical Mechanics.* New York: Springer, 1978 (translated from the Russian).

Arthur, W. Brian, Ermoliev, Yu. M., and Kaniovski, Yu. M. "Path Dependent Processes And The Emergence Of Macro Structure." *European Journal of Operations Research*, 1987, 30, 294-303.

Arthur, W. Brian. "Urban Systems and Historical Path Dependence," in Jess H. Ausubel and Robert Herman (eds.), *Cities and Their Vital Systems.* Washington, DC: National Academy Press, 1988, pp. 85-97.

Arthur, W. Brian. " 'Silicon Valley' Locational Clusters: When Do Increasing Returns Imply Monopoly?" *Mathematical Social Sciences*, 1990, 19, 235-251.

Axelrod, Robert. "A Landscape Theory of Aggregation." Institute of Public Policy Studies, University of Michigan Working Paper #332, prepared for delivery at the American Political Science Association annual meeting, Washington, D.C., August 29-September 1, 1991.

Besen, Stanley M. and Saloner, Garth. "The Economics of Telecommunications Standards." In R.W. Crandall and K. Flamm (eds.), *Changing the Rules: Technological Change, International Competition, and Regulation in Communication.* Washington: Brookings, 1989.

Burke, Steven. "New RTs aid competitiveness of IBM in workstation market." *PC Week*, August 1, 1988, 5(31), 90.

Chamberlin, Edward H. *The Theory of Monopolistic Competition.* Cambridge, MA: Harvard University Press, 1933.

Computer Technology Research Corp. *Technical Workstations In-Depth.* Holtsville, NY: Computer Technology Research Corp., May 1990a.

Computer Technology Research Corp. *Unix in the 1990s.* Holtsville, NY: Computer Technology Research Corp., July 1990b.

David, Paul. "CLIO and the Economics of QWERTY." *Papers and proceedings of the American Economic Association*, May 1985, 75, 332-336.

David, Paul, and Greenstein, Shane. "Selected Bibliography on the Economics of Compatibility Standards and Standardization." *Economics of Innovation and New Technology*, 1991, 1, 3-41.

Ethier, Wilfred. "Decreasing Costs in International Trade and Frank Graham's Argument for Protection." *Econometrica*, 1982, 50, 1243-1268.

Farrell, Joseph and Saloner, Garth. "Standardization, compatibility, and innovation." *Rand Journal of Economics*, Spring 1985, 16(1), 70-83.

Farrell, Joseph and Saloner, Garth. "Installed base and compatibility: Innovation, product preannouncements, and predation." *American Economic Review*, 1986, 76, 940-955.

Farrell, Joseph and Saloner, Garth. "Coordination through committees and markets." *Rand Journal of Economics*, Summer 1988, 19(2), 235-252.

Fraas, Arthur G., and Greer, Douglas F. "Market Structure and Price Collusion: An Empirical Analysis." *Journal of Industrial Economics*, September 1977, 26, 29-33.

Friedman, James W. *Oligopoly and the Theory of Games*. Amsterdam: North-Holland, 1977.

Fudenberg, Drew and Maskin, Eric. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." *Econometrica*, May 1986, 54, 533-554.

Gabel, H. Landis. "Open Standards in the European Computer Industry: The Case of X/OPEN." In H.L. Gabel (ed.), *Product Standardization and Competitive Strategy*. Amsterdam: North-Holland, 1987, 91-123.

Green, Edward and Porter, Robert. "Non-Cooperative Collusion Under Imperfect Price Information." *Econometrica*, 1984, 52, 87-100.

Hamel, Gary, Doz, Yves L. and Prahalad, C.K. "Collaborate With Your Competitors -- and Win." *Harvard Business Review*, July-August 1989, 67, 133-139, .

Hart, Sergiu and Kurz, Mordecai. "Endogenous Formation of Coalitions." *Econometrica*, July 1983, 51, 1047-1064.

Holland, John H. and Miller, John H. "Artificial adaptive agents in economic theory." *Papers and proceedings of the American Economic Association*, May 1991, 81, 365-370.

Hopfield, John J. and Tank, David W. "Computing with Neural Circuits: A Model." *Science*, August 8 1986, 233, 625-633.

Jorde, Thomas M., and Teece, David J. "Innovation And Cooperation: Implications For Competition And Antitrust." *Journal of Economic Perspectives*, Summer 1990, 4, 75-96.

Katz, Michael, and Shapiro, Carl. "Network Externalities, Competition, and Compatibility." *American Economic Review*, June 1985, 75, 424-440.

Katz, Michael, and Shapiro, Carl. "Technology Adoption in the Presence of Network Externalities." *Journal of Political Economy*, August 1986, 94(4), 822-841.

Kauffman, Stuart A. "Adaptation of Rugged Fitness Landscapes." in Daniel L. Stein (ed.). *Lectures on the Sciences of Complexity, volume 1*. Redwood City, CA: Addison-Wesley, 1989.

Krugman, Paul. "History Versus Expectations." *Quarterly Journal of Economics*, May 1991, 106, 651-667.

Krugman, Paul. "Trade, Accumulation, and Uneven Development." *Journal of Development Economics*, 1981, 8, 149-161.

Liebowitz, S.J. and Margolis, Stephen E. "The Fable Of The Keys." *The Journal of Law & Economics*, April 1990, 33, 1-25.

Nelson, Richard R. and Winter, Sydney G. *An Evolutionary Theory Of Economic Change*. Cambridge, MA: Harvard University Press, 1982.

Nicolis, Gregoire and Prigogine, Ilya. *Exploring Complexity, An Introduction*. New York: Freeman, 1989.

Owen, Guillermo. "Values of Games with a Priori Unions." In R. Hein and O. Moeschlin (eds.), *Essays in Mathematical Economics and Game Theory*. New York: Springer-Verlag, 1977, 76-88.

Panagariya, Arvind. "Increasing Returns, Dynamic Stability, and International Trade." *Journal of International Economics*, 1986, 20, 43-63.

Rajan, Roby. "Endogenous Coalition Formation in Cooperative Oligopolies." *International Economic Review*, 1989, 4, 863-876

Rubinstein, Ariel. "Equilibrium in Supergames with the Overtaking Criterion." *Journal of Economic Theory*, 1979, 21, 1-9.

Saloner, Garth. "Economic issues in computer interface standardization." *Economic Innovation and New Technology*, 1990, 1, 135-156.

Scherer, F. M. and Ross, David. *Industrial Market Structure and Economic Performance* (third edition). Boston: Houghton Mifflin Company, 1990.

Selton, Reinhard. "A Simple Model of Imperfect Competition Where Four Are Few and Six Are Many." *International Journal of Game Theory*, 1973, 2, 141-201.

Shapley, Lloyd S. and Shubik, Martin. "Pure Competition, Coalitional Power and Fair Division." *International Economic Review*, October 1969, 10, 337-362.

Sorensen, Kristiina. "Workstation Game Too Early To Call." *Digital Review*, July 11, 1988, 5(13), 106.

Stigler, George J. "A Theory of Oligopoly." *Journal of Political Economics*, 1964, 72, 44-61

Werden, Gregory J. and Baumann, Michael G. "A Simple Model of Imperfect Competition in Which Four are Few But Three are Not." *Journal of Industrial Economics*, March 1986, 34, 331-335.

Wright, Sewell. "The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution." *Proceedings of the International Congress of Genetics*, 1932, 1, 356-366.

FIGURE 1. A LANDSCAPE WITH TWO LOCAL OPTIMA

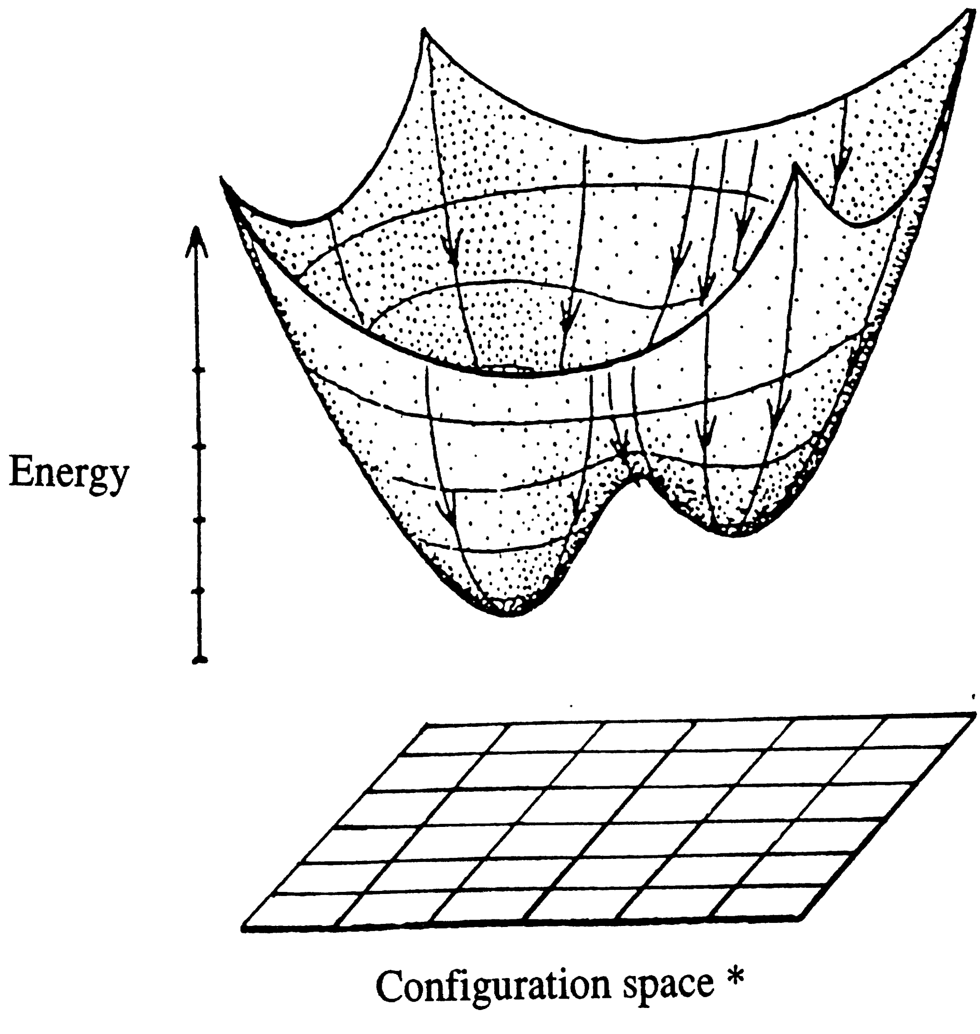

Adapted from Ralph Abraham and Christopher Shaw (1983, Part 2); used with permission of Ariel Press.

* The configuration space is an n-dimensional binary hypercube. The hypercube has one dimension for each firm indicating which of the two possible alliances that firm is in.

Table 1 -- The Two Configurations Predicted By Landscape Theory Analysis [a]

	Configuration 1	
	Alliance 1	*Alliance 2*
Configuration 2		
Alliance 1	Sun (28.9, S) Prime (1.0, S)	DEC (20.0, G) HP (11.5, G)
Alliance 2	AT&T (28.5, G) IBM (3.8, G)	Apollo (21.2, S) Intergraph (4.4, S) SGI (4.4, S)
Nearest empirical match [b]	Unix International	Open Software Foundation

[a] Size is shown in parentheses, along with whether the firm was a computer generalist (G) or technical workstation specialist (S). All firms had a design orientation.

[b] In configuration 1, only the IBM prediction is wrong.

FIGURE 1. A LANDSCAPE WITH TWO LOCAL OPTIMA

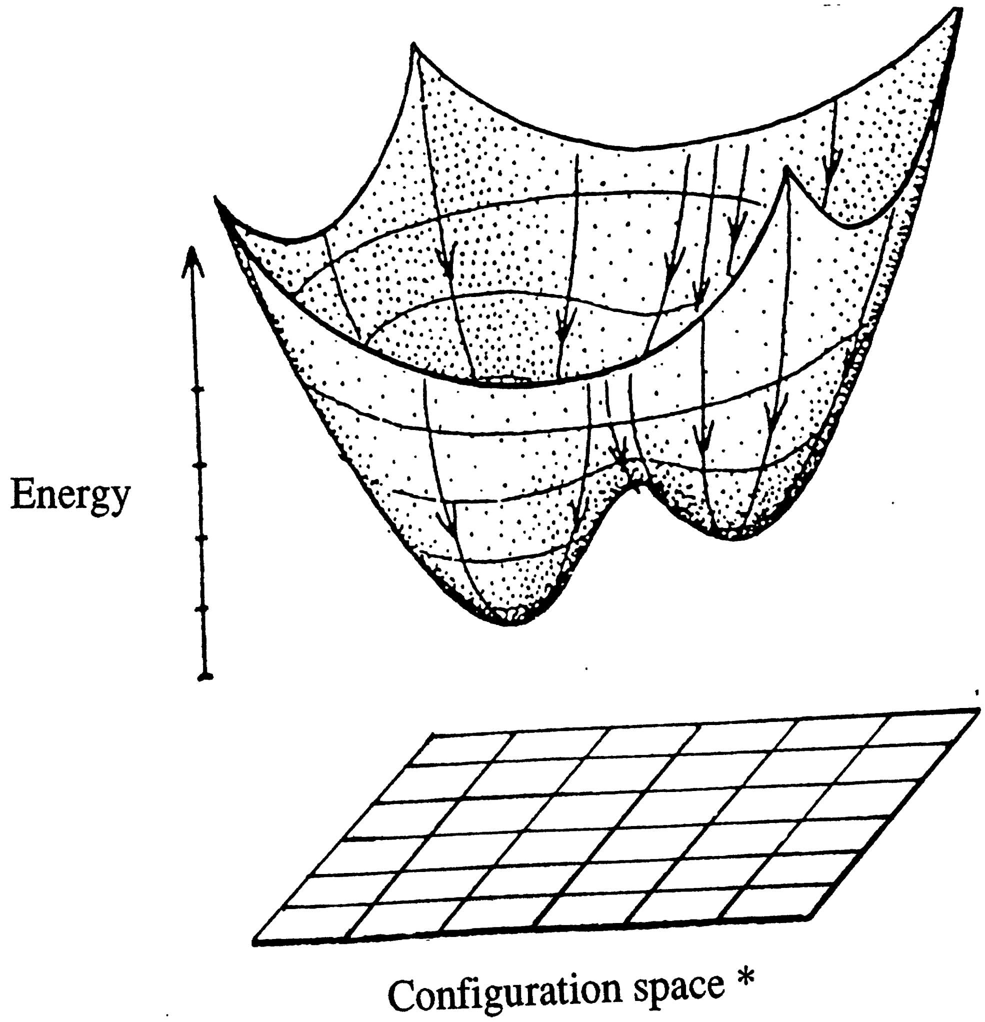

Adapted from Ralph Abraham and Christopher Shaw (1983, Part 2); used with permission of Ariel Press.

* The configuration space is an n-dimensional binary hypercube. The hypercube has one dimension for each firm indicating which of the two possible alliances that firm is in.