Positioning on a Multi-Attribute Landscape

Ron Adner[a], Felipe A. Csaszar[b], Peter B. Zemsky[c]

[a]Tuck School of Business, Dartmouth, 100 Tuck Hall, Hanover, NH 03755
Tel: (603) 646 9185, Fax: (603) 646 8711, Email: Ron.Adner@dartmouth.edu
[b]Ross School of Business, U. of Michigan, 701 Tappan Ave., Ann Arbor, MI 48109
Tel: (734) 615 4854, Fax: (734) 936 6631, Email: fcsaszar@umich.edu
[c]INSEAD, Boulevard de Constance, 77300 Fontainebleau, France
Tel: +33 1 60 98 34 08, Fax: +33 1 60 74 55 00, Email: peter.zemsky@insead.edu

August 27, 2011

Abstract

Understanding positioning is a central concern for strategy. We offer a rich but tractable formalization of competitive positioning that is explicit about how the success of firms' policy choices in the face of competition is affected by the multiple attributes via which firms can create value for consumers. On the supply side, our theory incorporates multiple organizational design choices; on the demand side, it incorporates heterogeneous buyers with preferences over multiple product attributes. Critical parameters are the extent of trade-offs that firms face when setting attribute levels and the degree of interactions among organization design decisions. We use a value-based approach to characterize competitive interactions in the marketplace. Three unexpected results emerge from bringing together the various elements of competitive positioning in a unified analytic framework. First, not all positions on the efficient frontier are viable. Second, in contrast to prior work on NK models of rugged landscapes, increases in business policy interdependence (i.e., increases in K) can decrease heterogeneity in viable positions. Third, market heterogeneity can be characterized in terms of either product attribute differentiation or business policy differentiation, and the relationship between the two is moderated by the extent of business policy interdependency.

Keywords: competitive positioning, value-based strategy, NK landscape

1 Introduction

Competitive positioning is a central concern of the strategy literature. For an individual firm, positioning entails a choice of how to compete in a given market. At an industry level of analysis, a key question is the extent to which a given environment supports heterogeneity in competitive positions and hence in the associated business policies (e.g., Nelson, 1991). In his influential work on the topic, Porter (1985, 1996) emphasizes that the existence of organizational trade-offs across multiple performance attributes (e.g., cost versus quality or ease of use versus feature richness) is what gives rise to the need for firms to make clear positioning choices. Yet despite being a central topic in the field of strategy, competitive positioning has rarely been the focus of analytic research. This gap is likely due to the breadth and integrative nature of the phenomena: a complete treatment of competitive positioning needs to incorporate organizational design in the presence of trade-offs, consumer demand in the presence of multiple performance attributes, and competitive interactions among firms in the marketplace. In this paper we seek to provide a parsimonious model of competitive positioning that captures all of these key elements.

Two of the most influential representations of firm positions are productivity frontiers and rugged landscapes. We incorporate both in our analysis. In a *productivity frontier* (Figure 1a), positions are represented as points in a two-dimensional space (see, e.g., Porter 1996:62; Saloner et al. 2001:61). Textbook depictions of the frontier show a smooth trade-off between the two dimensions. Positions inside the frontier are inefficient because both attributes can be improved by moving to the frontier. Porter (1996:61) asserts that moving to the frontier is an operational matter and that strategy is about the choice of a position on the frontier.

In a *rugged landscape* (Figure 1b), positions are represented as a vector of business policies that determine an overall fitness level (e.g., Levinthal, 1997). At the heart of this representation is the degree of interaction among discrete policy choices within the organization. The more interactions there are, the more rugged the landscape (i.e., the greater the number of local peaks). In the strategy literature, this representation has most frequently been used to study search by boundedly rational agents who can become trapped on local peaks. This representation offers a powerful way to formalize important policy interactions within organizations (e.g., Rivkin, 2000; Rivkin and Siggelkow, 2003).

Both approaches suffer from shortcomings. In the received NK literature in strategy, policy choices map onto a unidimensional performance measure (i.e., the fitness level). However, a satisfactory treatment of positioning requires that performance be considered in terms of multiple product attributes; this is what gives rise to the classic notion of trade-offs along a frontier. Conversely, while the frontier representation explicitly incorporates trade-offs among multiple product attributes, it is silent on the underlying business policy choices that give rise to positions. Moreover, both approaches lack explicit consideration of consumer choice and competitive interactions.[1]

We develop a novel approach to positioning that exploits the strengths of both the landscape and

[1]One notable exception is the work of Lenox, Rockart, and Lewin (2006, 2007), where firms on a rugged landscape engage in Cournot competition. However, the focus of these studies is still on a single performance attribute.

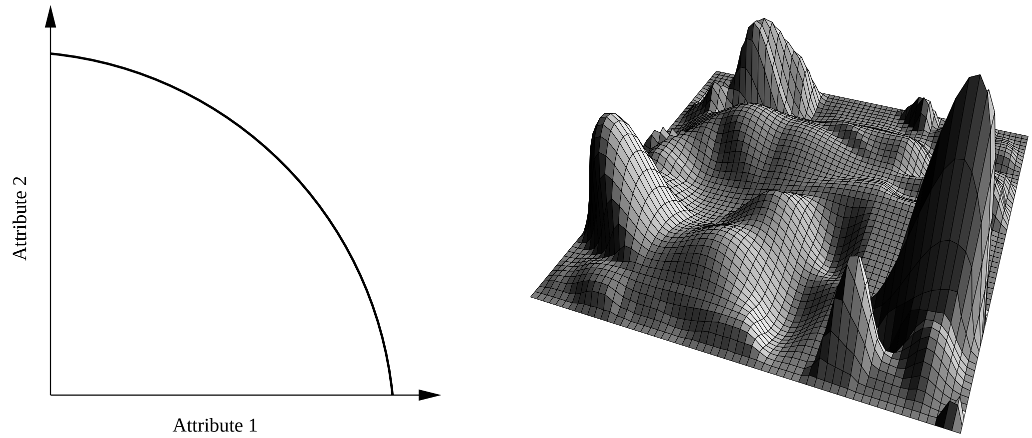

(a) A production frontier (b) A landscape

Figure 1: Examples of a production frontier and a landscape.

frontier representations. In addition, we incorporate consumer choice and competitive interactions using a value-based strategy approach (Brandenburger and Stuart, 1996). We characterize the extent of trade-offs among product attributes according to how changes in business policies that affect performance on one attribute also affect performance on other attributes. Thus we develop an explicit mapping between business policies and product attributes. We introduce a parameter that allows trade-offs among attributes to vary from negative (e.g., increasing the size of a car decreases its fuel efficiency) to zero (e.g., increasing car size does not affect color selection) to positive (e.g., increasing size increases safety).

We show that both the extent of interaction among business policies and the extent of trade-offs among product attributes have a fundamental impact on three aspects of industry heterogeneity: (i) the number of positions along the frontier that remain viable in the face of competition; (ii) the extent of heterogeneity in business policies among the viable positions; and (iii) the extent of heterogeneity among the products in the different viable positions. We also explore how these factors affect industry concentration and market shares.

Three unexpected results emerge from bringing together the various elements of competitive positioning in a unified analytic framework. First, not all positions on the efficient frontier are viable. Second, increases in business policy interdependence can decrease heterogeneity in viable positions. Third, market heterogeneity can be characterized in terms of either product attribute differentiation or business policy differentiation, and the relationship between the two is moderated by the extent of policy interdependency.

The paper proceeds as follows. Section 2 argues for why the study of positioning would benefit from an integrative modeling approach. Section 3 specifies our model, and Section 4 illustrates the mechanics of the model by working out an example. Section 5 presents the main results, and Section 6 concludes.

2 The Need for a New Approach

Although there is an extensive industrial organization (IO) literature on product differentiation that is relevant to competitive positioning, it is not adequate to the task of addressing the fundamental strategy questions that concern us here. For example, there is a vast literature on Hotelling models in which firms position themselves along a line or a circle and then compete. There is even some work that models, as we do, competition in a multi-attribute space with heterogeneous consumers (e.g., Lancaster, 1966, 1990; De Palma et al., 1985; Canoy and Peitz, 1997). However, one weakness of these approaches is the lack of a link to the underlying business policies that are required to occupy a given position. In fact, firms in IO models are typically assumed to face a given and smooth trade-off among positions (e.g., the Hotelling line) that is not grounded in any representation of the internal organization of the firm. This is limiting for strategy research on positioning, where a central concern is the viability of a firm's internal organization given the external competitive market environment.

The NK modeling methodology (originally developed in evolutionary biology; see Kauffman 1993) has been widely adopted in the strategy literature to model the effects of search by boundedly rational actors (see, e.g., Levinthal 1997; Rivkin 2000; Rivkin and Siggelkow 2007; Csaszar and Siggelkow 2010; and references therein). A distinctive advantage of this methodology is its ability to formally model organizational concepts (such as bounded rationality, modularity, organization design, and analogical thinking) that are central to the current understanding of strategy, but that were not amenable to rigorous analysis using traditional IO approaches.

A weakness of the NK methodology is that it has mostly developed as a line of thought that does not intersect with IO approaches. In particular, the great majority of NK models do not consider issues of competition, such as market share, concentration, and profits (Baumann and Siggelkow, 2010). A notable exception to this lack of integration is the work by Lenox, Rockart, and Lewin (2006, 2007), which, by combining an NK model with a Cournot model, has made novel predictions regarding the relationship between industry profits and the potential for interdependency among activities, as well as offered an alternative causal explanation for industry shakeouts.

A characteristic of all NK models used in the strategy literature is that they measure performance on a unidimensional scale called "fitness" (a name reminiscent of the NK model's biological origins). Although this translation from biology to strategy leads to a straightforward interpretation of fitness in terms of firm performance (e.g., profits), it prevents the model from being used in settings where performance is multidimensional. As we show in this paper, the assumption that performance is unidimensional has profound implications, as multidimensional spaces are fundamentally different from unidimensional spaces. For this reason, it is unclear how the predictions of the NK models in strategy would translate to settings where performance (and not just policy choices) is multidimensional.

For example, while in a unidimensional world, 2 is always above 1, in a multidimensional space it is unclear if, say, $(4, 2)$ is preferable over $(2, 4)$. An ordering relationship exists only when a point is superior to another in all of its dimensions (e.g., $(4, 4)$ is superior to or dominates $(2, 2)$). The

Entry

Firms enter with strategies denoted by the binary string $\mathbf{s}_i \in \{0,1\}^N$. There is a large number of entering firms, with at least one firm on each of the 2^N possible strategies.

Technology

Nature determines a multi-attribute landscape that maps each strategy onto performance along two dimensions as given by the functions $a_1(\mathbf{s}_i)$ and $a_2(\mathbf{s}_i)$.

Competition

Firms compete for buyers, who purchase the product that creates the most value for them according to the function $v(a_1(\mathbf{s}_i), a_2(\mathbf{s}_i); \alpha)$, where α varies across buyers and gives the relative weight placed on each of the two attributes.

Figure 2: Stages of the game.

fact that the ordering relationship is not well defined leads to sets of equivalent points.

This fundamental topological difference between unidimensional and multidimensional spaces has important implications for strategy, since many real-world competitive settings are indeed multidimensional. For example, the car industry competes in terms of prices, safety, mileage, etc., while the search engine industry competes in terms of comprehensiveness, usability, response time, etc. Moreover, it is precisely the lack of a well-defined ordering relationship in multidimensional spaces that makes the concept of positioning meaningful. More concretely, it is because a product with characteristics $(2, 4)$ is not clearly better than a product with characteristics $(4, 2)$ that it makes sense to ponder questions such as which one of the two to offer, or whether or not it makes sense to come up with product $(3, 3)$.

Addressing these questions requires a new approach that integrates multiple attributes, discrete policy choices, and heterogeneous demand. We present such an approach below.

3 Model

Our model has three stages, as summarized in Figure 2. In the first stage, a large number of firms enter the market with heterogeneous business strategies. Building on the NK model, we represent each firm's strategy as a bundle of business policies. Namely, a firm's strategy is represented as N binary business policy choices; we denote the *strategy* of a firm by $\mathbf{s} \in \mathcal{S}$, where $\mathcal{S} = \{0,1\}^N$. We interpret business policies as encompassing all the choices that affect the firm's performance, such as organization design and product design choices. In other words, a given $\mathbf{s}$ can be understood as a detailed description of the strategy of a firm, akin to the concept of a business model (Casadesus-Masanell and Ricart, 2010). We assume that there is entry of a large number of heterogeneous firms such that all 2^N possible strategies are represented in the industry.

In the second stage, the firms offer products whose performance on two key attributes a_1 and a_2 varies with a firm's strategy $\mathbf{s}$ according to the functions $a_1(\mathbf{s})$ and $a_2(\mathbf{s})$. The function $a_1(\mathbf{s})$ is modeled as a standard fitness function in an NK model (see Appendix A for a detailed description of the NK model). That it, there are N contribution functions, and each function depends on the value of $K+1 \leq N$ business policies. The function $a_2(\mathbf{s})$ is modeled similarly. We introduce trade-offs between the two attributes by allowing for correlation between the underlying contribution functions; this correlation determines the extent (if any) to which a high value of a_1 is associated with a high or low value of a_2.

We call the point $(a_1(\mathbf{s}), a_2(\mathbf{s}))$ the *position* of a firm with strategy $\mathbf{s}$. Note that a strategy corresponds to an N-dimensional point, while a position corresponds to a two-dimensional point. We believe that this mapping from $\mathbf{s}$ to a_i captures an essential characteristic of strategic management: that managers must deal with a large number of levers, whereas consumers care only about the smaller set of attributes that are visible to them. An important property of our assumptions, as in many business settings, is that managers do not have continuous dials to select their competitive position. Rather, specific competitive positions are determined by a series of underlying organizational and product design choices. For example, a hotel manager cannot directly control "customer comfort" but can control such parameters as bed size, decor, and staffing levels.

In the third stage, firms compete for customers who vary in the weight $\alpha \in [0, 1]$ that they place on the two product attributes. This is the source of demand heterogeneity. Customers purchase the product that creates the most value for them. Value created is given by the function $v(a_1(\mathbf{s}), a_2(\mathbf{s}); \alpha)$, which depends on the interplay of product attributes ($a_1(\mathbf{s})$ and $a_2(\mathbf{s})$) and customer preferences (α).

We say that a position is *efficient* if there is no position that dominates it on both attributes (i.e., it is on the production frontier).[2] We say that a position is *viable* if it has the highest value creation for at least one customer (i.e., the position is able to attract demand even when competing alternatives are located in all other positions).

It is important to highlight that our model implicitly incorporates a high level of competition, both in terms of entry and in terms of rivalry among the entrants. In the model this is given by the assumption of firms in every position and that competitive outcomes are tightly linked to value creation. We make these assumptions explicit in Appendix B, where we show that our results on viable strategies are equivalent to formal rational models that involve a large number of potential entrants, no fixed costs to entry, no capacity constraints, and perfect price discrimination.

The rest of this section details each part of the model.

3.1 Entry and Organizational Design

Our focus is not on game theoretic entry decisions by firms. Instead, we make the simple assumption that there is a large enough number of entrants into the industry such that all positions are occupied by at least one firm. We make this assumption for the sake of simplicity, but it is not necessary for our results. In Appendix B we show that other common, fully rational IO models can give rise

[2] Formally, efficiency requires that there does not exist an $\mathbf{s}' \in \mathcal{S}$ for which $a_1(\mathbf{s}') \geq a_1(\mathbf{s})$ and $a_1(\mathbf{s}') \geq a_1(\mathbf{s})$ with one inequality strict.

to the same results regarding efficient and viable positions as our simpler model. We show this for both a biform game (Brandenburger and Stuart, 2007) and for a traditional two-stage entry game (Tirole, 1988).

3.2 Technology and Trade-offs

We assume that what matters to consumers is the performance of the product they purchase (e.g., the durability and appearance of a roofing tile), rather than the business policies that give rise to the product (e.g., semi-automated versus fully automated production lines). We focus on the case in which there are two key performance attributes. Textbook examples of two-attribute settings are cost and quality in hotels (Saloner et al., 2001:61) and computing power and battery life in laptop computers (Spulber, 2004:218). We characterize the extent of trade-offs between the two attributes by using parameter ρ, which is defined below.

Recall that the level of attribute i is given by $a_i(\mathbf{s})$. In specifying each of these functions we follow the NK methodology (Levinthal, 1997). This methodology, widely adopted in the strategy literature (see, e.g., Rivkin, 2000; Gavetti et al., 2005; Rivkin and Siggelkow, 2007), creates random landscapes of a size determined by the parameter N and with a degree of interactions or "ruggedness" that is determined by the parameter K. The effect of K is to determine how nonlinear is the mapping between the business policy choices ($\mathbf{s}$) and the product attribute values (a_i). Mathematically, each attribute function is defined as

$$a_i(\mathbf{s}) = \frac{1}{N} \sum_{j=1}^{N} c_i^j(s_j; K \text{ other elements of } \mathbf{s}), \tag{1}$$

where $c_i^j(\cdot)$, called a "contribution function," determines the contribution of business policy j to attribute i as a function of the value of business policy j and the value of K other business policies. See Appendix A for a detailed description of the NK model and how contribution functions are set. Since each of the contribution functions can take values between 0 and 1 and since each attribute is an average of contribution functions, it follows that each attribute takes values between 0 and 1.

In many settings, the value of the attributes may exhibit trade-offs whereby increasing the level of one attribute requires decreasing the other. For example, the size of a car is usually negatively correlated to its fuel efficiency. Porter (1996:69) remarks that "trade-offs are essential to strategy. They create the need for choice and purposefully limit what a company offers." We can capture strong trade-offs between attribute levels by imposing that the contribution functions for the second attribute are perfectly negatively correlated with the contribution functions for the first attribute:

$$c_2^j(\cdot) = 1 - c_1^j(\cdot) \quad \text{for all } j,$$

so that $a_2(\mathbf{s}) = 1 - a_1(\mathbf{s})$.

Perfect negative correlation is a strong assumption, and in many settings one would expect that attribute levels would only be imperfectly correlated. Indeed, some attributes—such as the interior design and fuel efficiency of a car—might be largely independent. We introduce imperfect correlation by varying the number of contribution functions that are linked across attributes. Let $Q \leq N$ be the number of contribution functions for which $c_2^j(\cdot) = 1 - c_1^j(\cdot)$, and let the remaining contribution functions be independent.[3] The case of no trade-offs (and hence no correlation) between attribute levels corresponds to $Q = 0$.

In some settings, it is possible that attributes could be positively correlated. For example, if simpler product designs are easier to use and more reliable, then ease of use and reliability would be positively correlated. We can introduce positive correlation, in a similar fashion, by equating Q of the contribution functions, so that $c_2^j(\cdot) = c_1^j(\cdot)$, and allowing the rest to be independent. For $Q = N$ we have perfectly positive correlation (i.e., $a_2(\mathbf{s}) = a_1(\mathbf{s})$).

We define an overall measure of correlation $\rho = Q/N$ for the case of positive correlation and $\rho = -Q/N$ for the case of negative correlation. Thus, ρ can assume any value from -1 to $+1$. For $\rho = -1$ we have $a_2(\mathbf{s}) = 1 - a_1(\mathbf{s})$, for $\rho = 1$ we have $a_2(\mathbf{s}) = a_1(\mathbf{s})$, and for $\rho = 0$ we have that $a_1(\mathbf{s})$ and $a_2(\mathbf{s})$ are independent. The parameter ρ will be useful for presenting our results as it parameterizes the extent of trade-offs in an industry's technology.

3.3 Competition and Value Creation

We take a value-based approach to modeling market interactions. Such an approach starts with a precise statement of the set of actors in the industry and their value creation possibilities. The actors in our model are the set of entrants positioned across the competitive positions $\mathcal{S}$ and a finite number of buyers who vary in their preferences over the two attributes. We parameterize preferences with $\alpha \in [0,1]$ and denote the set of buyers by $\mathcal{A} \subset [0,1]$.

Each buyer has demand for one unit of the industry output. Value creation is increasing in both attributes, whose relative importance depends on consumer preferences. In particular, buyer α served by a firm with strategy $\mathbf{s} \in \mathcal{S}$ leads to a value creation of

$$v(\mathbf{s}, \alpha) = v_0 + \alpha \log(1 + a_1(\mathbf{s})) + (1 - \alpha) \log(1 + a_2(\mathbf{s})) - c, \tag{2}$$

where c is a constant marginal cost of production, v_0 is a constant in the consumer's willingness to pay (WTP), and the term $\alpha \log(1 + a_1(\mathbf{s})) + (1 - \alpha) \log(1 + a_2(\mathbf{s}))$ captures the effect of attributes and preferences on WTP. We simplify the analysis by assuming that $v_0 \geq c$, so that value creation is always positive.[4]

Because our primary interest in this paper is the set of viable positions in the industry rather

[3] Our modeling approach is inspired by Csaszar and Siggelkow (2010), who introduce positive association between contribution functions to capture the relatedness of the landscapes for organizations operating in different contexts. We have taken this approach to a multi-attribute setting and then introduced the possibility of negative association in order to capture trade-offs in the attributes.

[4] There is an analogous formulation in which attribute 2 reduces marginal costs instead of increasing WTP; in that case, the frontier maps the trade-off between cost and quality (e.g., Porter, 1996). In particular, a model with marginal costs given by $c - \log(1 + a_2(\mathbf{s}))$ and WTP given by $v_0 + \frac{\alpha}{1-\alpha} \log(1 + a_1(\mathbf{s}))$ has the same set of viable positions and market shares as the model studied in this paper.

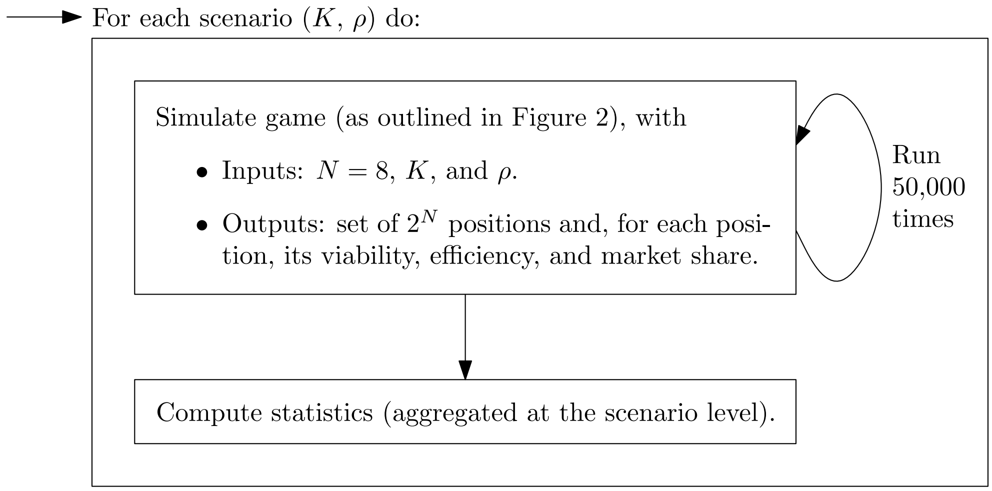

Figure 3: Simulation outline.

than the profits of individual firms, we do not specify a detailed model of competition. Instead , we assume that competitive rivalry is sufficiently intense to lead customers to be served by firms located at the position with the highest value creation for the specific customer. In short, a customer with preferences α is served by a firm positioned at $\mathbf{s}$ if and only if $v(\mathbf{s}; \alpha) \geq v(\mathbf{s}'; \alpha)$ for all $\mathbf{s}' \in \mathcal{S}$.

3.4 Simulation

To analyze the model, we set $N = 8$ and run a simulation over the entire parameter space defined by K and ρ. Namely, we vary K from 0 to 7 and vary ρ from -1 to 1 (in increments of 0.125, which given that $N = 8$, is the minimum increment allowed by the definition of ρ). This leads to 136 ($= 8 \times 17$) scenarios. To avoid interpreting results that are a function of a specific random draw, we run 50,000 simulations per scenario and report aggregate performance statistics at the scenario level (all the reported results are statistically significant at the 0.01 level at least). Each time a game is run, a new NK landscape is drawn. Figure 3 outlines the workflow used to run the games and to aggregate the simulation results.

In addition to using $N = 8$, we also run the simulation with other values of N (6 and 10). Consistently with the NK literature, in broad terms the results were not sensitive to the choice of N, but rather to the relative size of K with respect to N.

We assume that consumers are uniformly spread over the range $[0, 1]$. Because the model combines simulation (the NK landscape generation) with closed forms (the competition in terms of value creation), it is convenient to use discrete rather than continuous distributions; thus we assume that there are $M = 1000$ consumers equally spaced in the $[0, 1]$ range. If M is large enough then, for all practical purposes, the discrete approximation yields the same results as a continuous specification. Consistent with this, we find that results are essentially equivalent in robustness tests where we set $M = 500$ and $M = 2000$.

3.5 Relation to Industrial Organization Models

Our model incorporates many important features from IO, especially the extensive literature on product differentiation (see for example, Tirole 1988:ch. 7 and Vives 1999:ch. 6). For example, a common way to model demand in the IO literature on differentiation is to have a set of customers who vary in their preferences and who each buy up to one unit of output from among those offered by the firms in the industry. This is precisely our approach. In general, the demand side of our model draws on many elements of prior work. However, our model is distinctive in the way it uses an explicit assumption of multiple product attributes to blend horizontal and vertical differentiation.

Although there are notable exceptions (discussed below), IO models of differentiation often represent a firm's position by a single product attribute, which we will here refer to generically as a. In models of horizontal differentiation, shifts in a increase the WTP of some customers and reduce that of others. For example, in Hotelling's classic linear city model of horizontal differentiation, a is simply the position on the line and customers vary in their ideal point on the line. Therefore, a marginal increase in a increases the WTP of customers to the right of the current position and decreases the WTP of customers to the left.

In the case of vertical differentiation, value is increasing in a for all customers.[5] In these models, customers vary not in their ideal point but rather in their WTP for the attribute. This is usually modeled as a multiplicative term αa, where α varies across customers and parameterizes the WTP for the attribute. An early analysis of competition under vertical differentiation is Shaked and Sutton (1982). The demand side of Shaked and Sutton (1982) is a special case of our model with $a_2 = 0$.

In introducing two distinct product attributes a_1 and a_2 we build on the work of Lancaster (1966), who emphasizes that differentiation in a market is often not uni-dimensional. Lancaster (1966) and subsequent work such as Ben Akiva et al. (1989) follow Hotelling and study a generalized form of horizontal differentiation in which customers vary in their ideal level on each attribute. In contrast, we build on the vertical differentiation literature and assume that WTP is monotonically increasing in both attributes.

Our model incorporates vertical differentiation in order to capture the emphasis in the strategy literature on trade-offs between two value-creating attributes. We also incorporate horizontal differentiation by allowing the relative weight placed on a_1 and a_2 to vary across customers. Recall that in our model WTP is given by Equation (2). Consider two products, one with attributes (a_1, a_2) and the other with attributes (a_1', a_2'). If $a_1 > a_1'$ and $a_2 > a_2'$, then all customers have a higher WTP for the first product. Yet if $a_1 > a_1'$ but $a_2' > a_2$, then some customers (depending on their value of α) will have a higher WTP for the first product and others for the second—this corresponds to horizontal differentiation.

While the demand side of our model closely follows the existing IO literature on differentiation, the supply side is a more significant departure. In terms of technology, it is common in IO to assume that firms choose their positions directly and face smooth exogenous trade-offs. For example, the

[5] Examples of vertical attributes are computer speed and transportation safety.

firms in Hotelling models choose their location a; as discussed previously, this choice smoothly increases the WTP of some customers while smoothly reducing it for others based on an exogenous loss function. Motta (1993) adds cost asymmetries to Shaked and Sutton (1982) by assuming that costs are a smooth increasing function of quality. This intellectual tradition is strongly consistent with the emphasis on trade-offs and a smooth frontier in Porter (1996) and Saloner et al. (2001). In contrast, we assume that firms do not directly choose their position but rather their business policies. The frontier in our model is not necessarily smooth but rather the result of complex interactions among the set of business policies.

The IO literature does introduce stochastic elements into models of differentiation. However, this is done on the demand side. For example, De Palma et al. (1987) take a standard Hotelling model and add an independent random shock to the value that a customer gets from consuming the product of a given firm. Similarly, Rhee (1996) takes Shaked and Sutton (1982) and adds an independent random shock to the value that a customer gets from consuming the product of a given firm.[6] The introduction of these error terms is motivated verbally by the existence of additional unobserved product attributes and heterogeneous consumer preferences, but these are not formally modeled. While this has proven a productive approach for IO, we believe that there are advantages of our more structured approach using the NK methodology, at least for strategy research. In particular, we identify parameters of the underlying stochastic process—namely, K and ρ—that have meaningful managerial interpretations as policy interdependence and attribute correlation, respectively.

Finally, as discussed extensively in Appendix B, we take a different approach to competitive interactions. Most IO models of differentiation involve price competition in which firms simultaneously commit to a single price for all customers. This often leads to involved game theoretic analysis. The most natural interpretation of our model within the IO tradition is that firms engage in price competition but with perfect price discrimination, so that firms are essentially competing independently for each customer. We feel our assumption has empirical validity in some markets and has the advantage of a much simpler, more transparent analytic structure.

4 Understanding Positioning on a Multi-Attribute Landscape

In this section we illustrate our analysis of positioning on a multi-attribute landscape. We start with a given set of possible positions and then identify the efficient and viable positions. For each of the viable positions, we identify their market share. These mechanics are important for understanding the paper's main results, which characterize patterns in these variables over a large number of such landscapes.

We illustrate the analysis for the multi-attribute landscape given in Figure 4a, which is for the case of $N = 3$. In this case there are $2^3 = 8$ possible positions, and the figure shows the level of

[6]The stochastic element of consumer WTP serves to lessen competitive pressures because firms with similar positions are less perfect substitutes. As a result, the pressure to differentiate is reduced. The equilibrium level of differentiation, a topic first raised in Hotelling (1929), is perhaps the central question in this literature.

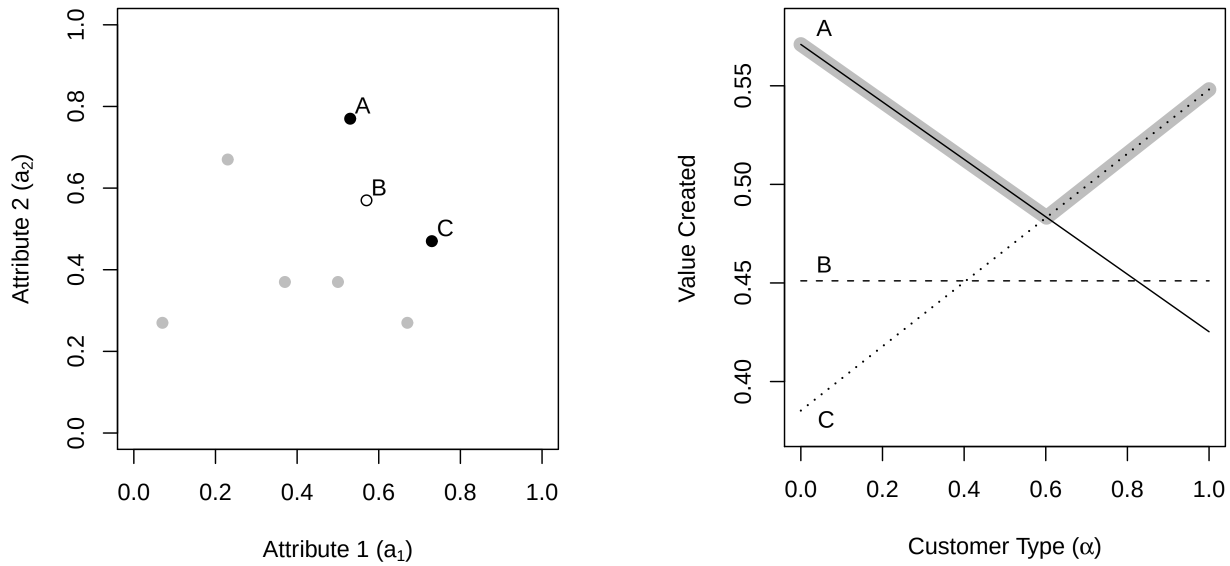

(a) Example of a frontier

(b) Value created by products on the frontier

Figure 4: Example of a frontier and how it relates to consumer valuations. In panel (a), the dots represent viable (●), efficient but not viable (○), and dominated (●) positions.

attribute 1 and attribute 2 for each of the positions.[7]

Value creation is increasing in both attributes. In Figure 4a, the points A, B, and C are the set of efficient positions and constitute the efficient frontier in this example (i.e., these are the only three points in the plot for which there are no points located to their "northeast"). All remaining points are inefficient because they offer less of attribute 1 and attribute 2 than one of the points on the frontier and therefore create less value. Firms located at these inefficient points never have maximum value creation for a customer once the positions on the frontier are occupied.

An efficient position is not necessarily viable. Viability requires that a position has the highest value creation for at least one customer (i.e., the position is able to attract demand even when competing alternatives are located in all other positions). A position with only marginally more of attribute 1 or attribute 2 than other positions might be efficient but need not be viable. To identify the viable positions, Figure 4b plots the value creation of the three efficient points. Because value creation varies across buyers, the horizontal axis represents the consumer type, α. A position has added value if there are some buyer types for which it has the highest value creation.

[7]Figure 4 was created using $\rho = 0$ and $K = 1$ as well as the following contribution functions for attribute 1:
$c_1^1(s_1{=}0, s_2{=}0){=}0.9, c_1^1(s_1{=}0, s_2{=}1){=}0.2, c_1^1(s_1{=}1, s_2{=}0){=}0.0, c_1^1(s_1{=}1, s_2{=}1){=}0.7,$
$c_1^2(s_2{=}0, s_3{=}0){=}0.2, c_1^2(s_2{=}0, s_3{=}1){=}0.1, c_1^2(s_2{=}1, s_3{=}0){=}0.7, c_1^2(s_2{=}1, s_3{=}1){=}0.4,$
$c_1^3(s_3{=}0, s_2{=}0){=}0.5, c_1^3(s_3{=}0, s_2{=}1){=}0.8, c_1^3(s_3{=}1, s_2{=}0){=}0.1, c_1^3(s_3{=}1, s_2{=}1){=}0.9,$
and the following contribution functions for attribute 2:
$c_2^1(s_1{=}0, s_3{=}0){=}0.6, c_2^1(s_1{=}0, s_3{=}1){=}0.1, c_2^1(s_1{=}1, s_3{=}0){=}0.5, c_2^1(s_1{=}1, s_3{=}1){=}0.2,$
$c_2^2(s_2{=}0, s_3{=}0){=}0.8, c_2^2(s_2{=}0, s_3{=}1){=}0.5, c_2^2(s_2{=}1, s_3{=}0){=}0.2, c_2^2(s_2{=}1, s_3{=}1){=}0.5,$
$c_2^3(s_3{=}0, s_1{=}0){=}0.9, c_2^3(s_3{=}0, s_1{=}1){=}0.7, c_2^3(s_3{=}1, s_1{=}0){=}0.5, c_2^3(s_3{=}1, s_1{=}1){=}0.1;$
which produced the following attribute values (as a function of s_1, s_2, s_3):
$a_1(0{,}0{,}0){=}0.53, a_1(0{,}0{,}1){=}0.37, a_1(0{,}1{,}0){=}0.57, a_1(0{,}1{,}1){=}0.50, a_1(1{,}0{,}0){=}0.23, a_1(1{,}0{,}1){=}0.07, a_1(1{,}1{,}0){=}0.73, a_1(1{,}1{,}1){=}0.67,$
$a_2(0{,}0{,}0){=}0.77, a_2(0{,}0{,}1){=}0.37, a_2(0{,}1{,}0){=}0.57, a_2(0{,}1{,}1){=}0.37, a_2(1{,}0{,}0){=}0.67, a_2(1{,}0{,}1){=}0.27, a_2(1{,}1{,}0){=}0.47, a_2(1{,}1{,}1){=}0.27.$

An efficient position is not necessarily viable. Viability requires that a position has added value so that a firm positioned there can successfully compete for buyers against firms in other positions. A position with only marginally more of attribute 1 or attribute 2 than other positions might be efficient but need not be viable. To identify the viable positions, Figure 4b plots the value creation of the three efficient points. Because value creation varies across buyers, the horizontal axis represents the consumer type, α. A position has added value if there are some buyer types for which it has the highest value creation.

Positions A and C are viable, but not position B. Position A has added value for customers with $\alpha < 0.6$ and position C has added value for customers with $\alpha > 0.6$, so firms can occupy these positions and capture demand. In contrast, position B does not have added value for any buyer type: either A or B create more value for each buyer type. For this reason, despite being on the frontier, Position B is not viable and cannot be occupied profitably in the face of rivals in Positions A and C.[8]

The market share of each viable position corresponds to the proportion of buyers for which it has added value. If buyers are uniformly distributed over the range $[0, 1]$, then the market share of position A is 60% and of position C is 40%.

5 Results

We use our model to explore how attribute correlation (ρ) and policy interdependence (K) affect key strategic variables in a multi-attribute world. We start by characterizing the topography of the competitive landscape and then proceed to examine the implications of this topography for the relative market shares associated with different positions. We conclude by examining how the interaction of ρ and K shapes the extent of heterogeneity in a market.

We begin our analysis with a characterization of the 2^N possible positions that exist in the industry. We identify (i) inefficient positions that lie within the frontier, (ii) efficient positions that lie on the frontier, and (iii) viable positions on the frontier that can attract customers even when rivals are positioned elsewhere on the frontier. In doing so, we are explicit in accounting for how the presence of rivals affects competitive outcomes.

5.1 Efficient Positions

As shown in Figure 5, we find overall that the expected number of efficient positions in an industry is decreasing in both attribute correlation (ρ) and policy interdependence (K). If attributes are perfectly correlated ($\rho = 1$), such that there is effectively only one attribute, there is only one position on the frontier. This is because, regardless of customer preferences (α), every position that offers less value than the single maximum is inefficient, or dominated by this single maximum. In contrast, if attribute trade-offs are extreme ($\rho = -1$) and so increasing the performance of one

[8] It is certainly possible for intermediate positions like B to be viable; for example, if the B line in Figure 4b were shifted upward by 0.1 then B would have added value for consumers with intermediate levels of α.

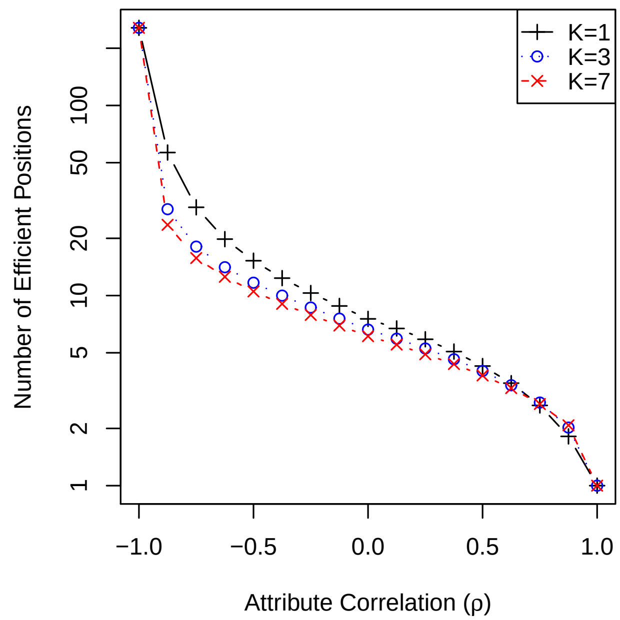

Figure 5: Number of efficient positions as a function of attribute correlation ρ and policy interdependence K.

attribute requires a corresponding decrease in performance of the other attribute, then each of the 256 possible positions lies along the efficiency frontier.[9] Note that, as attribute correlation is reduced from $+1$, we depart from the single-attribute landscape found in standard NK models. Conversely, as the correlation is increased from -1, we depart from the strongest form of trade-offs.

Figure 6 offers a graphical intuition for the effect of ρ on the number of efficient positions. Recall from Section 4 that the efficient points are those for which no points lie to their northeast. If $\rho = 1$ and so all positions lie along the line that extends upward from the origin at a 45-degree angle, then the only efficient point is the single one farthest from the origin. In contrast, if $\rho = -1$ then the positions extend downward at a -45-degree angle and so there is no single point positioned northeast of another. As ρ values approach 0 from both directions, such that there is less and less correlation between the attributes, we find a shift from the linear distributions at the extremes towards an increasingly nebulous distribution. Points that lie within this "cloud" are dominated by a few points on its upper right edge. These points on the upper right edge define the efficiency frontier (these are the solid and hollow black points in Figure 6).

Figure 5 also shows that the number of efficient positions is decreasing in policy interdependence K.[10] This is a surprising result because it seems to contradict the well-established result in the

[9] There are $256 (= 2^8)$ positions because $N = 8$ in the simulations and each policy can take one of two values.

[10] A careful examination of Figure 5 reveals that, at $\rho = 0.75$ and $\rho = 0.875$, the ordering of the three lines is slightly altered (i.e., the $K = 1$ line appears just below the other two lines). This nonmonotonicity with respect to the effect of K is due to a geometrical property that only arises in high ρ settings. To illustrate this phenomenon, imagine a cloud with $\rho = 1$ and $K = 0$ (i.e., a perfectly straight and upward-sloping collection of points, as for every $\mathbf{s}$, $a_1(\mathbf{s}) = a_2(\mathbf{s})$). If ρ is decreased to 0.875 (i.e., so that the two fitness functions differ by a single contribution function), then: (i) the entire set of possible positions is distributed along two parallel lines (because, by Equation (1), $a_2(\mathbf{s})$ must now be equal to either $a_1(\mathbf{s}) - c_1^j(0) + c_2^j(0)$ or $a_1(\mathbf{s}) - c_1^j(1) + c_2^j(1)$, where $c_i^j(\cdot)$ is the one contribution

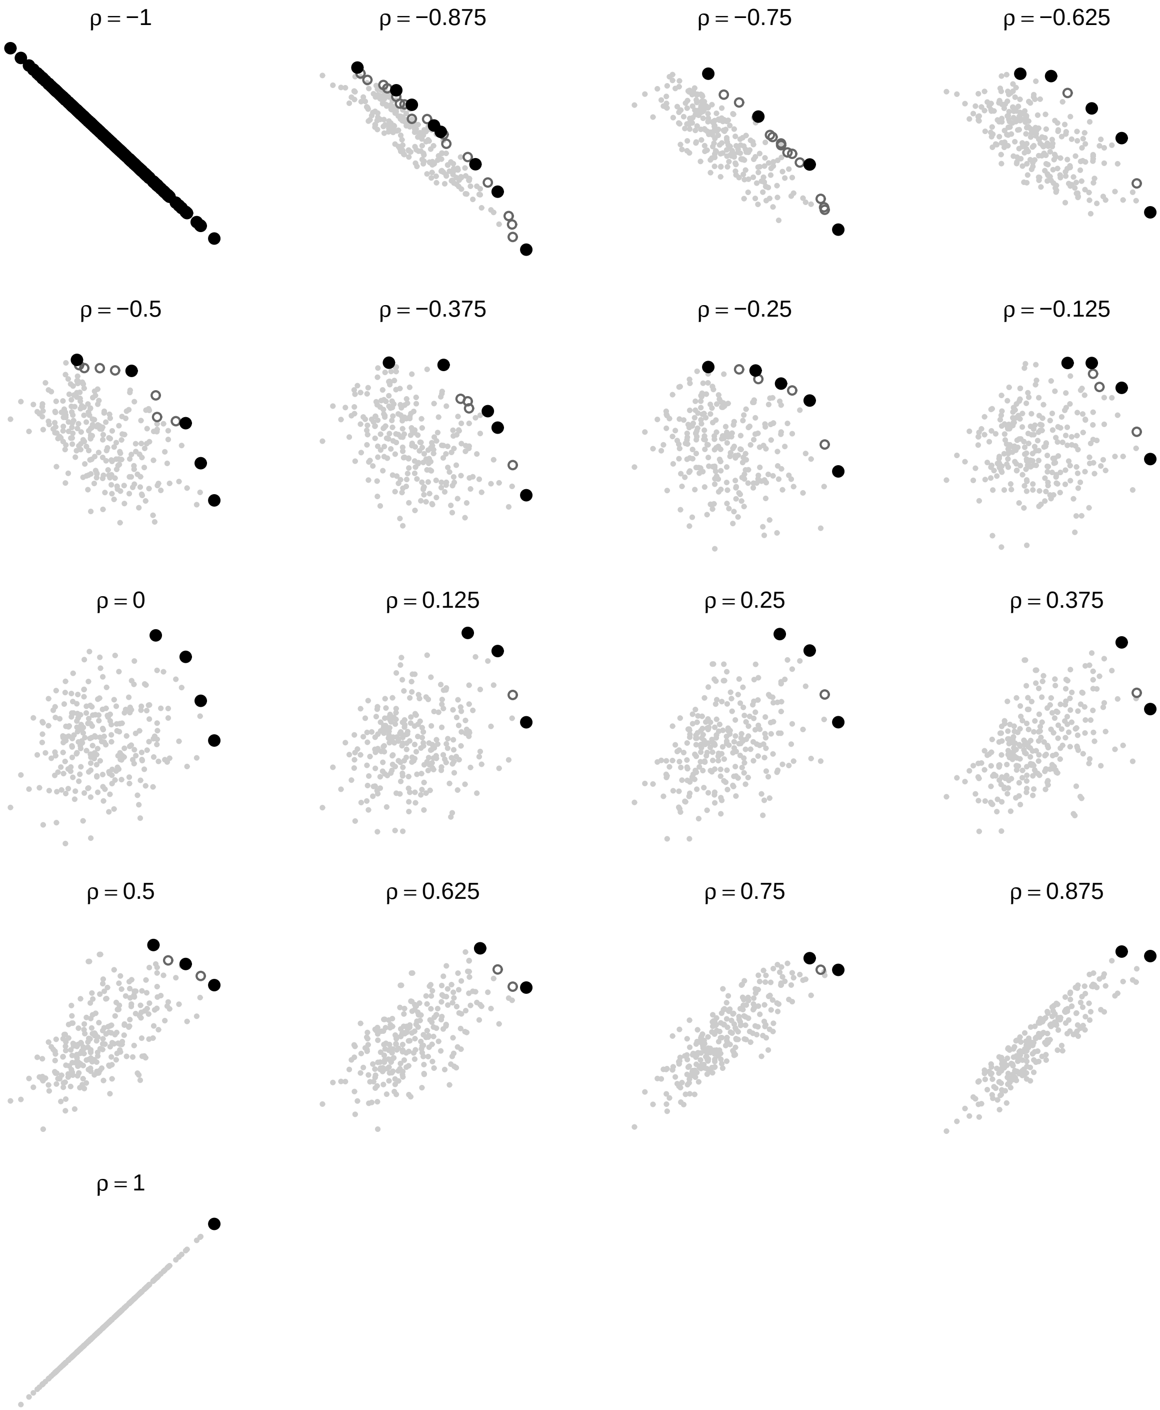

Figure 6: Graphical intuition of the effect of attribute correlation ρ on the number of viable (•), efficient but not viable (○), and dominated (●) positions. Each panel plots a single simulation, and its x and y axes are (respectively) a_1 and a_2; in all panels, $N = 8$, $K = 3$, and ρ varies from -1 to 1.

K	Mean	Standard Deviation	Skewness	Kurtosis
0	0.500	0.071	0.000	2.440
1	0.500	0.087	0.000	2.616
2	0.500	0.095	0.000	2.737
3	0.500	0.099	0.000	2.785
4	0.500	0.100	0.000	2.815
5	0.500	0.101	0.000	2.827
6	0.500	0.102	0.000	2.832
7	0.500	0.102	0.000	2.832

Table 1: Descriptive statistics of the fitness levels of landscapes (averaged statistics computed from 50,000 simulations per level of K and with $N = 8$).

NK literature, that the number of local peaks is increasing with K (Levinthal, 1997). To resolve this contradiction, it is important to note that there is not a one-to-one mapping between concepts in the single-attribute NK model and the multi-attribute landscape that we examine here. Most relevant for the current result is the fact that a local peak is not the same as an efficient point. To understand what drives our result, it is useful to look at the descriptive statistics of NK landscapes (see Table 1). The key observation to be made from this table is that, as K increases, the kurtosis of the landscape also increases. Recall that kurtosis is a measure of the extent to which the variance in a distribution is due to infrequent extreme deviations rather than frequent small deviations (Barnett and Lewis, 1994). As kurtosis increases with K, the absolute number of outliers decreases but the magnitude of their deviation increases. In the context of multidimensional landscapes, where efficient points are by definition outliers, this implies that the number of efficient points is decreasing in K.

5.2 Viable Positions

The viability of a position is determined by its ability to attract customer demand even if the other efficient positions are occupied by rivals. The number of viable positions is of interest because it can be interpreted as capturing an important aspect of industry heterogeneity.

Figure 7 plots the number of viable positions as a function of attribute correlation (ρ) for different levels of policy interdependence (K). As was the case with efficient positions, the likelihood of "standing out from the cloud" is decreasing in both attribute correlation (ρ) and policy interdependence (K).

At either $\rho = 1$ or $\rho = -1$, the number of viable positions is equal to the number of efficient positions. This is explained using the same arguments as for the number of efficient positions.

function that differs between $a_1(\mathbf{s})$ and $a_2(\mathbf{s})$); and (ii) it can be shown that, when $K = 0$, there is a 50% probability that the two parallel lines will lead to just one efficient point. With higher values of K, the distribution of the set of possible positions looks less and less like a pair of parallel lines and more and more like a cloud. However, the less well-ordered distribution leads to a greater chance of producing more than one efficient point. This K-driven effect is evident only at extremely high values of ρ, since lower ρ values lead to a cloudlike distribution for all values of K.

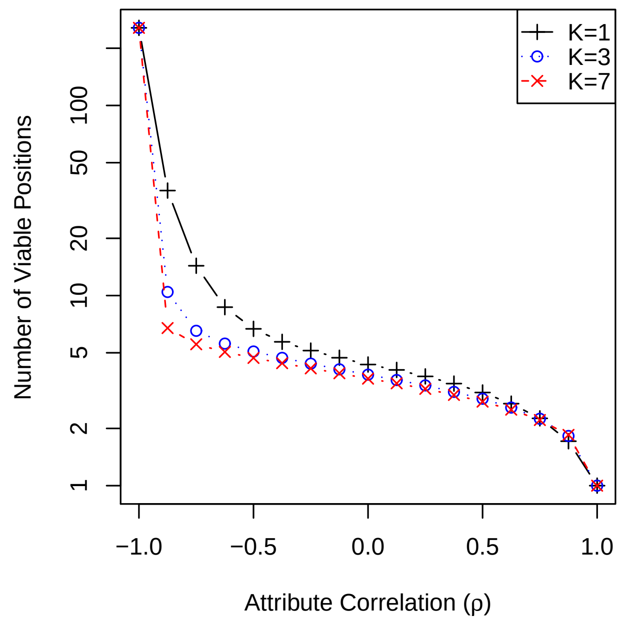

Figure 7: Number of viable positions as a function of attribute correlation ρ and policy interdependence K.

That is, if $\rho = 1$ then there is only one efficient (and thus viable) position. If $\rho = -1$, then all positions are efficient and lie along a perfectly straight, downward-sloping line and—because the log utility function is convex (which assures that all the positions along this downward-sloping straight frontier will maximize the utility of at least one customer)—all these positions, too, are viable.

For intermediate values of ρ, the qualitative trends regarding correlation and interdependence are the same as for the case of efficient positions, but the number of viable positions is substantially lower than in that case. This result is expected because the conditions for being viable are more stringent than those for being efficient. In particular, a position that is not close enough to the upper right edge of the cloud will not be able to capture customers, who will be better served by other, better located rivals, as illustrated in Figure 6.

By clearly identifying the existence of viable positions, we follow Porter (1996) in showing that there are important positioning choices that must be made by firms even after they have reached the efficiency frontier. An overarching regularity is that the extent of trade-offs increases the number of positions that can be supported in the landscape (i.e., the more negative is ρ, the more viable positions there are). In practical terms this suggests that, when trade-offs are high, choosing a positioning becomes a more elaborate decision because there are more viable positions from which to choose.

5.3 Market Shares

So far we have shown that not all positions are efficient and that not all efficient positions are viable. Next we show that not all viable positions are equivalent in their ability to generate demand. We

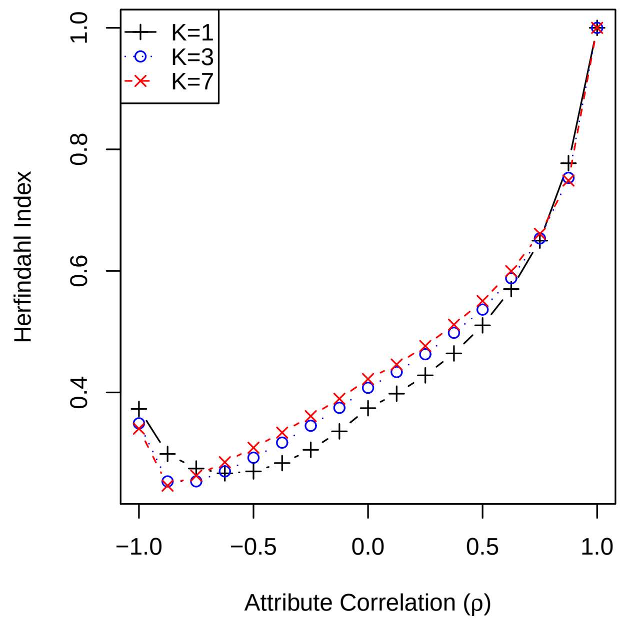

Figure 8: Herfindahl index as a function of attribute correlation ρ and policy interdependence K.

start by characterizing market concentration via a Herfindahl index across positions.[11] Figure 8 shows the relationship between market concentration as a function of attribute correlation (ρ) and policy interdependence (K). The figure indicates that concentration generally increases with ρ. This is consistent with the decrease in viable positions observed with increasing ρ, as reported in Figure 7: fewer viable positions lead to greater market shares. Similarly, we find that concentration increases with K, which is consistent with the reduction in viable positions that was found to accompany an increase in policy interdependence (shown in Figure 7).[12] Only at the extremely low levels of $\rho \leq 0.75$ does concentration increase slightly. The reason is that, when ρ is strongly negative, the numerous viable positions in between the extremes of the cloud are almost perfect substitutes for its neighbors, pushing the market share of these points close to zero, in effect leaving the bulk of the market to be supplied by either of the two extreme positions.

To better gauge heterogeneity in demand across positions, Figure 9 explores what we call the "market shares at the extremes"—that is, the sum of the market shares of the first viable position at the upper left of the efficient frontier and the last viable position at the lower right of the efficient frontier. Interestingly, we find significant heterogeneity across positions. Whereas the Herfindahl index in Figure 8 mostly increases with attribute correlation ρ, Figure 9 shows a richer relationship: the positions that lie at the extremes of the frontier capture a disproportionate share of sales throughout the range, which results in a U-shaped profile. At $\rho = 1$, there is only one viable position and so it holds the entire market. At $\rho = -1$, all 256 positions are viable but the two positions at the extreme ends of the frontier capture well over half the market.

The logic driving this large market share of the extremes is that such positions have no com-

[11]Formally, the Herfindahl index we compute is the sum of the squared market share of each of the viable positions.
[12]The slight alteration in the line ordering at $\rho = 0.875$ is explained by the arguments presented in footnote 10.

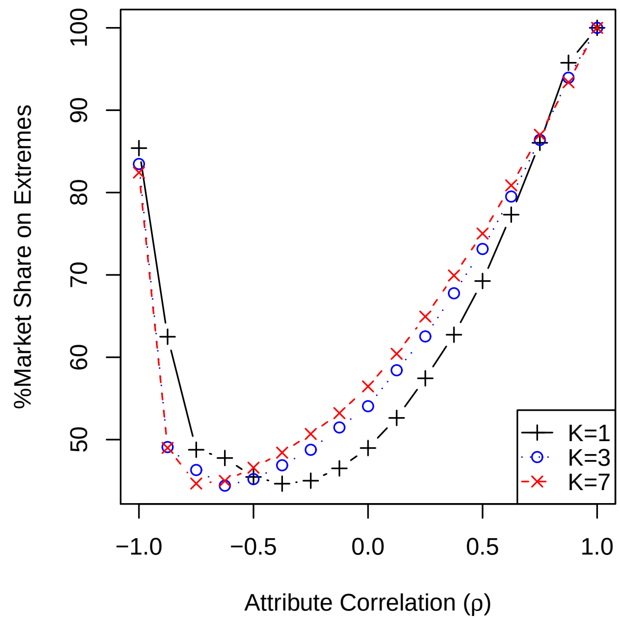

Figure 9: Fraction of market share served by the two viable positions at the extremes of the frontier, as a function of attribute correlation ρ and policy interdependence K.

petitors on one side (i.e., the upper left extreme has no competitors to the left and vice versa for the lower right position) and so—given that our model assumes consumers to be uniformly distributed ($\alpha \sim U[0,1]$)—there is a large segment of customers who are better served by the extreme positions. In other words, the extreme positions are the best possible choice for customers with a high WTP for one of the attributes. A real-world example is the highest-quality producer of high-fidelity equipment being preferred by audiophiles, who seek even better sound than what is available. Note, however, that this result is sensitive to the distribution of customers. For instance, if the bulk of consumers are located in the middle of the market (and so can be characterized by, e.g., a low-dispersion Normal distribution centered around $\alpha = 0.5$), then a firm would find positioning at the extremes to be less attractive; the converse would be true if the bulk of customers had extreme preferences.

5.4 Policy and Attribute Differentiation

We now consider two critical dimensions of heterogeneity: *product* heterogeneity, which we measure in terms of differences in attribute levels (i.e., differences between (a_1, a_2) pairs); and *business policy* heterogeneity, which we measure in terms of differences in policy choices (i.e., differences between $\mathbf{s}$'s). To gauge both types of heterogeneity, we create measures of the distance between the two positions at the extremes and consider how changes in correlation (ρ) and interdependence (K) affect these distances. Greater distance reflects greater heterogeneity in the population.

Figure 10 plots *product heterogeneity* measured as the Euclidean distance in attribute space between the extremes. In the context of Figure 4a, this equates to the length of a straight line

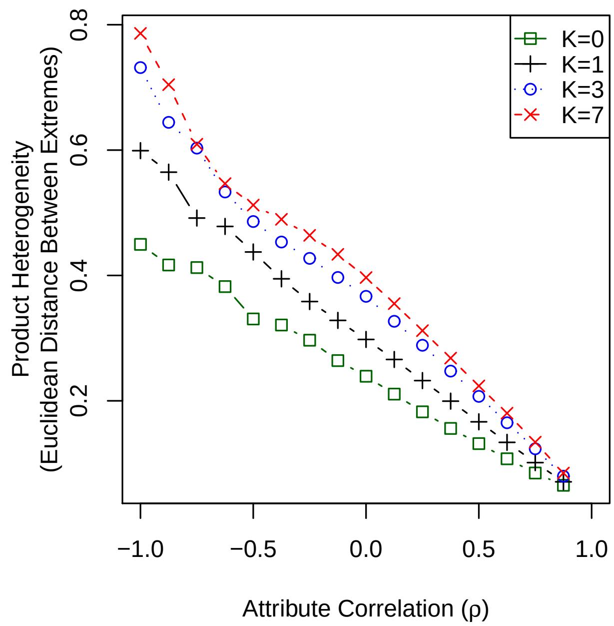

Figure 10: Distances in *attribute* space between the positions in the extremes as a function of attribute correlation ρ and policy interdependence K.

connecting points A and C. Increases in ρ result in a reduction in the Euclidean distance between the extremes.[13] This corresponds with the effect of ρ illustrated in Figure 6: the viable frontier, and the distance between viable points, shrinks as ρ approaches unity. In contrast, increases in K increase the Euclidean distance. This is because, as documented in Table 1, increases in K result in more pronounced outliers (i.e., higher kurtosis).

To measure *business policy heterogeneity* we use the Hamming distance between the policy choices of the positions at the extremes of the frontier. The Hamming distance is the number of bits that differ between two sets of policy choices. For example, the Hamming distance between 11000011 and 11000000 is 2, and the Hamming distance between 11111111 and 00000000 is 8. For presentation clarity, in our figures we show normalized Hamming distances (i.e., we divide the Hamming distance by 8). Whereas distance in attribute space can be thought of as reflecting differences in the physical manifestations of products, distance in policy space is analogous to differences in the underlying organizational structure, strategy, and routines that give rise to the product itself.

Figure 11 plots the normalized Hamming distance between the strategies of the two viable positions located at the opposite extremes of the frontier as a function of ρ and K. This figure shows a number of interesting relationships that connect attribute space with policy space. To understand these relationships, it is helpful to bear in mind that our model maps strategies to attributes and attributes to consumer valuations. In shorthand notation, we say that the model goes from $\mathbf{s}$-space into a-space into v-space. Now we look at how distances and positions are

[13]We do not plot values at $\rho = 1$, since in that case there is only one viable position and thus no distance between extremes.

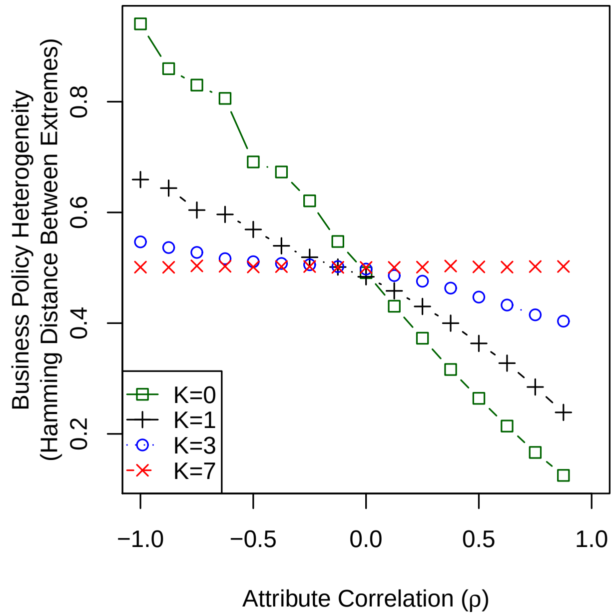

Figure 11: Distances in *policy* space between the positions in the extremes as a function of attribute correlation ρ and policy interdependence K.

transformed when moving from one space to another: in a-space (e.g., Figures 4a and 6) the distance between points indicates the extent of similarity between products (i.e., as points get closer together, the products get more alike); and the relative position of points indicates their relative value (i.e., points to the northeast are more valuable than points to the southwest). Thus there is a direct relationship between a-space and v-space. In contrast, the relationships between $\mathbf{s}$-space and a-space and between $\mathbf{s}$-space and v-space are not direct; instead they depend nontrivially on the value of K.

When $K = 0$, flipping a single bit of the strategy $\mathbf{s}$ from 1 to 0 affects the contribution of just that one bit: a small (one-bit) change results in a small change in attribute performance. In a low-K environment there is usually a direct (positive) relationship between the number of bits that are switched and the magnitude of the change in performance. Therefore, when K is low, Hamming distance in policy space behaves similarly to Euclidean distance in attribute space (i.e., the lines for $K = \{0, 1, 3\}$ in Figure 11 are all downward sloping, just as in Figure 10). But when K is high (e.g., $K = 7$), flipping a bit from 1 to 0 affects not only the contribution of that one bit (either up or down) but also the contribution of K other bits (either up or down). This means that a small change in $\mathbf{s}$ can result in either a large or a small change in performance. In a high K environment, there is a high level of variation in the effect of changing one policy on performance; in other words, the relationship becomes increasingly random as K increases. Thus, when $K = 7$, changing one bit (which changes the contribution of all the bits) is equivalent to changing all eight bits. In Figure 11 this is reflected in the horizontal line for $K = 7$, which results directly from the randomness of fitness shifts in response to policy changes.

Some interesting, and potentially testable, implications stem from comparing figures 10 and

11. First, policy interdependence K plays a much more important role in driving qualitative differences in policy heterogeneity than in attribute heterogeneity (i.e., shifts in K have a much more pronounced impact on the slopes in Figure 11 than in Figure 10). Second, whereas attribute heterogeneity always increases with K, policy heterogeneity only increases with K when attribute correlation is positive, but decreases with K when attribute correlation is negative. This mediation effect of policy interdependence (K) happens because, as mentioned above, as K increases the relationship between distances in $\mathbf{s}$-space and a-space becomes more random.

Finally, the previous analyses suggest that policy interdependence (K) plays an important role in determining the equifinality of different strategies. For example, if policy interdependence is low (e.g., $K = 0$) and there are strong trade-offs between the attributes (e.g., $\rho = -1$), then firms in the extremes of the efficient frontier will probably use quite different strategies (e.g., 00000000 versus 11111111) and will offer products that are perceived to be quite different (e.g., $(0.1, 0.9)$ and $(0.9, 0.1)$; perhaps this is the case of McDonald's versus Starbucks in the restaurant industry). But if policy interdependence is high (e.g., $K = 7$) and there are synergies between the attributes (e.g., $\rho = 0.75$), then different strategies could lead to similar products (e.g., Toyota's and Ford's approaches to delivering a hybrid vehicle may be quite different in policy space but can still result in cars that look quite similar in attribute space).

5.5 Robustness Check for Consumer Preferences

Apart from the robustness considerations and checks mentioned in Section 3.4, we also analyzed the model under an alternative specification in which we assume consumer preferences have linear trade-offs ($v(\mathbf{s}, \alpha) = v_0 + \alpha a_1(\mathbf{s}) + (1 - \alpha)a_2(\mathbf{s}) - c$), rather than the log formulation. The results overall are highly consistent with those derived from the main specification. The sole qualitative difference is that, at extremely negative values of ρ, the linear case supports a lower number of viable positions than does the log case. The reason is that, under linear preferences, if $\rho = -1$ (and thus the cloud of points form a left-to-right downward diagonal) then the two points in the extremes capture all the customers. This does not happen under log preferences, as these are more convex than linear preferences, and thus enable positions which lie farther away from the extremes of the cloud to capture market share. The predictions of the two alternative specifications become consistently more similar as ρ increases.

6 Discussion

This paper presents a parsimonious theory of positioning that considers both the interdependence of business policies, as articulated in the NK modeling literature, as well as an explicit treatment of competition and demand as developed in the IO literature. We examine how these factors interact in the context of a market in which consumers assess value creation along two different attribute dimensions.

Our investigation extends the literature on models of positioning in two ways. First, it departs from the assumption, implicit in IO models, that firms can directly choose what position to occupy in the attribute space. Second, it departs from the assumption of NK models, that fitness landscapes have just one dimension of performance. By combining elements of NK and IO models into one, we leverage the strengths of each approach. Additionally, our model considers both consumer heterogeneity and product market competition. This allows us to be explicit about how a rival's position in one part of the market affects outcomes in other parts. This addresses the call of Baumann and Siggelkow (2010) for analyzing how competition affects the landscape that firms must navigate.

Although each of these departures is straightforward, together they give rise to a landscape whose topography is significantly different from that characterized by traditional, single-attribute approaches. Using this structure, we can formally address a number of important strategy questions, including "What drives the number of viable positions in an industry?" and "How does viability differ from efficiency?" By decoupling business policies from product attributes, our novel use of a multidimensional NK structure allows us also to ask a rich set of questions regarding industry heterogeneity—for example, "What is the extent of heterogeneity among firms in an industry in terms of positions, business policies, and market shares?" We show that answers to such questions depend nontrivially on the joint impact of trade-offs between the product attributes (ρ) and the interdependence among business policy choices (K).

We identify three unexpected results. First, not all positions on the efficient frontier are viable. Second, in contrast to prior work on NK models of rugged landscapes, increases in business policy interdependence (i.e., increases in K) can decrease heterogeneity in viable positions. In particular, we qualify the strategy literature on NK models, by showing that results derived in a single-attribute setting do not necessarily generalize beyond that setting. Third, market heterogeneity can be characterized in terms of either product attribute differentiation or business policy differentiation, and the relationship between the two is moderated by the extent of business policy interdependency.

When comparing our work to the traditional NK models in strategy, we show that introducing even just one additional attribute opens up a multiplicity of viable positions in the market. While a casual analysis might have supposed that shifting from one attribute to two would allow the number of viable positions to double from one to two, we show that, depending on the extent of attribute correlation (ρ), the number of viable positions can increase dramatically.

An interesting area of application of our model is the research on business models and competitive advantage (e.g., Casadesus-Masanell and Ricart, 2010; Zott and Amit, 2008). By positing a clear set of mechanisms and deriving various relationships—among (for example) the Hamming distance, the extent of attribute trade-offs, and the extent of policy interdependence—our model presents testable hypotheses regarding the range of business model heterogeneity in the face of competition in a market.

There a numerous avenues for extending our model and analysis. Our focus has been on the extent and nature of heterogeneity, yet other important links between competitive positioning and firm performance in a multi-attribute setting remain to be explored. For instance, further work could relax some of our assumptions and study settings with fewer entrants, higher costs, or fixed pricing. In terms of methods, subsequent work could move toward the game-theoretic IO literature or toward the boundedly rational search of traditional evolutionary NK models. We consciously sought to avoid such elements in this initial contribution but nonetheless see them as promising directions for future work.

There are at least two important ways in which game-theoretic interactions could be incorporated into our model. As shown formally in Appendix B, in IO terms our model features free entry and perfect price discrimination. One avenue to enrich the strategic interactions among firms is to introduce a fixed cost of entry and extend the definition of viability to include having sufficient customer demand to cover the fixed costs. It is natural that there could then be multiple equilibrium configurations of viable positions and that firm and industry profits would vary depending on the particular equilibrium on which firms were able to coordinate. One could then study how the model's parameters, including now the level of fixed costs, affects the importance of managing equilibrium selection for performance. A second avenue towards a more classic IO analysis is to have a few firms locate on the landscape and then quote a single price for all customers. Firm positioning choices will then tradeoff the imperative to increase value creation against the desire to move apart to lessen price competition. Although much more complex to analyze, such a setting would allow for some firms to choose inefficient positions.

For strategy research, we find the possibility to introduce boundedly rational search into the model might be at least as promising as the more involved game theoretic analysis. Much of the strategy research conducted using NK models has examined how different modes of organizing local search yield different outcomes in the context of a single-attribute rugged landscape. It would be fruitful to explore how the shift from a single-attribute to a multi-attribute landscape affects competitive outcomes when firms are myopic and how different search heuristics perform at uncovering viable positions. Some results could be quite different from those derived using previous NK models, because in our model the fitness (i.e., market share and profits) of a strategy is not fixed but rather depends on where the other firms locate.

Given that positioning is such a central concern for both strategy practitioners and researchers, we believe that furthering our understanding of its determinants and implications is worthwhile. Our paper advances this agenda by developing a formal theory of positioning that brings together the notions of trade-offs, interdependencies, heterogeneous consumer demand, multiple value attributes, and competitive interactions in a single coherent model. An important by-product of taking a formal approach is that we are able to provide unambiguous definitions of important constructs such as efficiency, trade-offs, viability, and even positioning itself. Furthermore, we are able to ground key notions that have become part of the strategy lexicon (e.g., the "shape" of the technology frontier) by deriving them from basic economic principles, and use these to posit clear mechanisms (i.e.,"how" explanations) connecting the inputs and outputs of the process of positioning. We hope that this approach can serve as a platform for future investigations of key questions in competitive strategy.

A Appendix: Description of the NK Model

The NK landscape methodology allows modeling the performance of a general class of systems. Although originally developed to model biological systems (Kauffman, 1993), the NK methodology has been used extensively to model firms (e.g., Levinthal, 1997) and products (e.g., Sommer and Loch, 2004). Arguably, this methodology has been appealing to modelers of organizations, as it provides novel ways to formally analyze core organizational issues such as bounded rationality, modularity, interactions, and organizational search.

The NK landscape methodology encompasses a family of models. This appendix explains how to compute a standard NK landscape and gives pointers to papers that have extended the standard model.

An NK landscape is a function that maps the state of a system onto a measure of its performance, which is customarily called *fitness*. The system is assumed to have N components, and each component can exist in a number of states. For example, imagine a portable computer made of $N = 3$ components—screen, battery, and CPU—and suppose that each component can exist in one of two states: the screen can be small or large, the battery can be low-capacity or high-capacity, and the CPU can be slow or fast.

The contribution of a given component to the fitness of the system depends not only on the state of that component but also on the states of the other components with which it interacts. In the case of our computer: the screen's contribution to fitness depends on its size and on CPU speed (e.g., the combination of a large screen and a slow CPU yields especially sluggish performance); the battery's contribution to fitness depends on its capacity and on screen size (e.g., batteries are more rapidly depleted by larger screens); and the CPU's contribution to fitness depends on its speed and on the battery's capacity (e.g., a fast CPU works poorly with a weak battery).

The components and interactions in an NK model can be equivalently represented as a directed graph or an interaction matrix. For instance, the computer example corresponds to

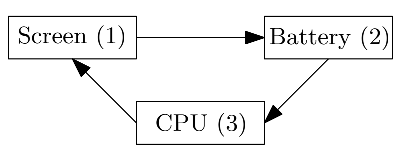

or $\mathbf{I} = \begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix}$.

In the *directed graph* representation, each box denotes a component and each arrow an interaction. In the *interaction matrix* representation, cell (i, j) contains a 1 if component i depends on component j (i.e., if the graph representation includes an arrow from j to i). Because the fitness of a component depends on its own state, the diagonal of the matrix is filled with ones.

How much each component contributes to the product's overall fitness is described by the N contribution functions. The *contribution function* i maps the states of the components that affect the contribution of component i onto the fitness contribution of component i. Continuing with our computer example, since the screen's contribution to fitness depends on screen size and CPU speed and since each component can be in two different states, it follows that the screen's contribution function can take one of four values: $c_{\text{screen}}(\text{small}, \text{slow})$, $c_{\text{screen}}(\text{small}, \text{fast})$, $c_{\text{screen}}(\text{large}, \text{slow})$, and $c_{\text{screen}}(\text{large}, \text{fast})$.

The standard NK model assumes that each component can be in either of two states (0 and 1), so each contribution function i can take $2^{(\#\text{ of ones in row } i \text{ of } \mathbf{I})}$ different values. That model also assumes that each component depends on K other components (i.e., that the interaction matrix has $K + 1$ ones per row). Thus, the ongoing example has $K = 1$ as each component depends on one other component apart from itself. Parameter K controls the degree of interrelatedness among the system's different components. With higher K, it is more likely that changing the state of one component will have an effect on the contribution of other components. This parameter can be seen as a way to "tune" the *complexity* that underlies the mapping between choices and fitness.

The fitness of a given position is defined as the sum of the fitness contributions for that position (normalized by N so that fitness values will be comparable across landscapes of different N's). Namely,

$$f(\mathbf{s}) = \frac{1}{N} \sum_{i=1}^{N} c_i(s_i; K \text{ other elements of } \mathbf{s}),$$

where $\mathbf{s} = (s_1, \ldots, s_N)$ represents the state of each component of the system.

To illustrate how fitness is computed, suppose the contribution functions of the three components in our example are as follows:

Contribution of screen

	screen	CPU	
$c_1($	0	, 0	$) = 0.9$
$c_1($	0	, 1	$) = 0.5$
$c_1($	1	, 0	$) = 0.3$
$c_1($	1	, 1	$) = 0.7$

Contribution of battery

	battery	screen	
$c_2($	0	, 0	$) = 0.8$
$c_2($	0	, 1	$) = 0.0$
$c_2($	1	, 0	$) = 0.0$
$c_2($	1	, 1	$) = 0.3$

Contribution of CPU

	CPU	battery	
$c_3($	0	, 0	$) = 0.1$
$c_3($	0	, 1	$) = 0.5$
$c_3($	1	, 0	$) = 0.2$
$c_3($	1	, 1	$) = 0.9$

Thus, the fitness of a portable computer with <u>small screen</u>, <u>weak battery</u>, and <u>fast CPU</u> would be:

$$f(0, 0, 1) = \frac{1}{3}\left(c_1(0, 1) + c_2(0, 0) + c_3(1, 0)\right) = \frac{1}{3}(0.5 + 0.8 + 0.2) = 0.5.$$

It is sometimes useful to represent NK landscapes as hypercubes, where each node represents a position and each link a connection between neighboring states (states that differ by one element). The hypercube representation of the landscape for our example appears in Figure 12a.

Although NK landscapes are N-dimensional objects, a useful imagery is to think about them as three-dimensional surfaces, where the dimensions lying parallel to the floor represent policy decisions, and the vertical dimension represents fitness. Because as K increases a landscape has more local peaks, the equivalent of K in the three-dimensional imagery is the ruggedness of a surface. Figures 12b and 12c illustrate low- and high-K landscapes.

The main property of NK landscapes is that, the more interactions a system has (i.e., the higher is K), the less similar are the fitness values of neighboring positions. If $K = 0$ then the fitness landscape represents a smooth surface whose global maximum can be found by hill climbing; as K

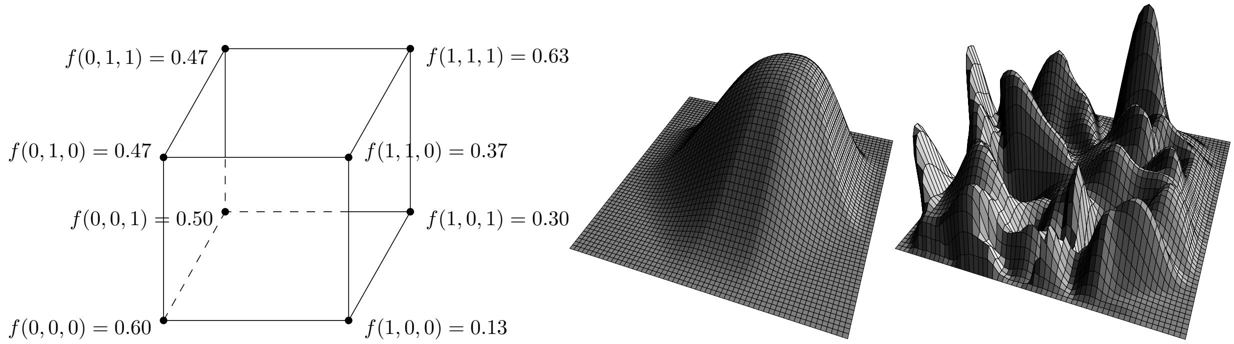

(a) Hypercube representation (b) Low K surface (c) High K surface

Figure 12: Graphical representations of NK landscapes.

increases, the landscape becomes more rugged or multipeaked and thus, hill climbing is unlikely to find the global maximum; and if K is maximal ($K = N - 1$) then the landscape represents an extremely rugged or spiky surface on which there is no correlation between the fitness of neighboring points. Bounded rationality in this context is usually conceptualized as a search process that is not omniscient about the whole landscape. Because few analytic results are known about NK landscapes (for one of the few exceptions, see Durrett and Limic, 2003), these models are typically analyzed via simulation.

Some of the management phenomena explored using models derived from the standard NK model are the search for dominant designs (Levinthal, 1997), imitation and replication of strategies (Rivkin, 2000; Csaszar and Siggelkow, 2010), coordination among organizational units (Rivkin and Siggelkow, 2003; Siggelkow and Levinthal, 2003), and the use of analogies by firms facing novel environments (Gavetti et al., 2005).

B Appendix: Equivalency of the Model with Bertrand and Biform Variations

In this appendix we show that the set of viable strategies is equivalent to the set of equilibrium strategies of two fully rational models of competition: (i) a biform game (Brandenburger and Stuart, 2007) that uses coalitional games to model competition (MacDonald and Ryall, 2004) and (ii) a two-stage, noncooperative game that is common in the IO literature with an assumption of price discrimination.

B.1 Model Variations

The biform and two-stage variations share a number of features. First, in both there are a large number (at least 2^N) of profit maximizing potential entrants. Second, both variations have a two-stage structure. In the first stage, the potential entrants decide whether to enter and, if they do so, they choose a strategy. As in the base model, an entrant's strategy is an element of $\mathcal{S} = \{0, 1\}^N$. We denote the number of firms that actually enter in stage I by m and their strategy choices by $\mathcal{E} = \{\mathbf{s}_1, \ldots, \mathbf{s}_m\}$. Third, value creation is exactly as in the base model. That is, an entrant's strategy determines the attribute levels of the firm's product according to the functions $a_1(\mathbf{s})$ and $a_2(\mathbf{s})$. We assume that these attribute functions are common knowledge. There is a constant marginal cost of production given by c, and firms compete for a set of customers $\mathcal{A} \subset [0, 1]$ whose WTP is given by $v(a_1, a_2; \alpha)$.

The two model variations differ in how competition takes place in stage II. In the classic two-stage approach, we assume that there is Bertrand competition with price discrimination. That is, each firm simultaneously names a price for each customer; the price of firm i for customer α is denoted by $p_{i\alpha}$. Customers buy one unit of the product that gives them the most surplus. As is typical under Bertrand competition, if customers are indifferent between two or more products, they are assumed to buy the one that generates the most profit for the selling firm. The profit of firm i in the two-stage game is then given by

$$\Pi_i^{\mathrm{TS}} = \sum_{\alpha \in \mathcal{A}} (p_{ia} - c) I_{i\alpha},$$

where $I_{i\alpha}$ indicates whether firm i sells to customer α.

In the biform game, we solve for the core of the coalitional game involving the m entrants and the customers in $\mathcal{A}$. Denote by $[\underline{x}_i, \overline{x}_i]$ the range of payoffs to firm i in the core. Then the profit of firm i in the biform game is given by $\Pi_i^{\mathrm{B}} = \gamma \underline{x}_i + (1 - \gamma) \overline{x}_i$, where γ is a parameter reflecting the confidence of firms in their ability to negotiate with customers within the constraints provided by the core (Brandenburger and Stuart, 2007).

We solve the two-stage game for a pure strategy subgame perfect equilibrium. That is, we first solve for the Nash equilibrium of the Bertrand pricing game in stage II for any set of entrants. We then require that entry decisions in stage I form a Nash equilibrium with profit functions given by

Π_i^{TS} from equilibrium play in stage II. Similarly, for the biform game, we solve for Nash equilibrium entry strategies in the first stage, where payoffs are now given by Π_i^{B}. As is common with entry models, we restrict our attention to equilibria in which all entrants have strictly positive profits in equilibrium, for both variations.

B.2 Characterization of Stage II Payoffs

As is standard, we work backwards and begin by characterizing the stage II payoffs for both variations. We start with the biform game by defining value creation and added value in our setting. Let the function $V(\mathcal{E}, \mathcal{A})$ be the total value creation (equivalently, economic surplus) for any set of entrants $\mathcal{E}$ and any set of customers $\mathcal{A}$. Given constant marginal costs that are the same for all firms, total value creation is maximized when each customer buys the product for which she has the highest WTP. Hence, we can decompose total value creation as follows:

$$V(\mathcal{E}, \mathcal{A}) = \sum_{a \in \mathcal{A}} \max_{\mathbf{s} \in \mathcal{E}} v(a_1(\mathbf{s}), a_2(\mathbf{s}); \alpha) - c.$$

The added value of entrant i for any set of entrants $\mathcal{E}$ is defined as the increase in value creation from including i in the game:

$$AV_i(\mathcal{E}) = V(\mathcal{E}, \mathcal{A}) - V(\mathcal{E} \setminus \mathbf{s}_i, \mathcal{A}).$$

We know from Brandenburger and Stuart (2007) that added value places an upper bound on the payoff of a player in a coalitional game. Further, Chatain and Zemsky (2007) identify a class of models with buyers and suppliers and constant marginal costs of production, in which the core always exists and payoffs are proportional to added value. It is straightforward to verify that our biform game satisfies the conditions in Chatain and Zemsky (2007) (specifically, the assumptions A1 and A2 required for their Proposition 1). Hence we have the following result.

Lemma B.1 *In the biform model, the profits of each entrant are proportional to its added value:*

$$\Pi_i^B(\mathcal{E}) = \gamma AV_i(\mathcal{E}).$$

We turn now to the classic two-stage variation. Given the assumptions of price discrimination and constant marginal costs, the second stage of this model can be decomposed into independent Bertrand price games where entrants compete in prices for each customer in $\mathcal{A}$. Each customer is then served by the entrant with the greatest value creation for that customer and the firm's profit is the increase in value creation over the next highest value creation. This is precisely the added value for that customer. Summing over all customers now yields the following result.

Lemma B.2 *In the two-stage variation, the profits of each entrant are given by its added value:*

$$\Pi_i^{TS}(\mathcal{E}) = AV_i(\mathcal{E}).$$

Proof Consider the equilibrium prices for a given customer $\alpha \in \mathcal{A}$. In a standard Bertrand price game, firms sell a homogeneous product but vary in their marginal costs. In equilibrium, the lowest-cost firm sells with a margin equal to its cost advantage while the firm with the next lowest cost prices at its marginal cost. While firms in our model have the same marginal costs but vary in the WTP of customers, this has little substantive impact on the equilibrium analysis (Vives, 1999:ch. 5). The firm with the second highest WTP prices at marginal cost and the firm with the highest WTP is then able to charge a margin exactly equal to its advantage in WTP. Letting $v_{i\alpha} = v(a_1(\mathbf{s}_i), a_2(\mathbf{s}_i); \alpha)$, an equilibrium price vector is given by[14]

$$p_{i\alpha} = \begin{cases} c & \text{if } \max\limits_{j \in \{1,\ldots,m\} \setminus i} v_{j\alpha} \geq v_{i\alpha}, \\ c + v_{i\alpha} - \max\limits_{j \in \{1,\ldots,m\} \setminus i} v_{j\alpha} & \text{otherwise.} \end{cases}$$

Thus, an entrant's profit is

$$\begin{aligned} \Pi_i^{\mathrm{TS}}(\mathcal{E}) &= \sum_{\alpha \in \mathcal{A}} (p_{i\alpha} - c) \\ &= \sum_{\alpha \in \mathcal{A}} \max\{0, v_{i\alpha} - \max_{j \in \{1,\ldots m\} \setminus i} v_{j\alpha}\} \\ &= V(\mathcal{E}, \mathcal{A}) - V(\mathcal{E} \setminus \mathbf{s}_i, \mathcal{A}) \\ &= AV_i(\mathcal{E}). \quad \blacksquare \end{aligned}$$

Comparing these two lemmas reveals that the stage II profit function of the classic two-stage game is a special case of the profit function in the biform game with $\gamma = 1$.

B.3 Equilibrium Characterizations

Let $\mathcal{V}^{\mathrm{B}} \subseteq \mathcal{S}$ be the set of strategies used in the equilibrium of the biform game, and let $\mathcal{V}^{\mathrm{TS}} \subseteq \mathcal{S}$ be the set of strategies used in the equilibrium of the two-stage game. Recall that $\mathcal{V}$ is the set of viable strategies as defined in Section 3.

Entry in stage I depends only on whether or not an entrant expects positive profits in stage II. Given that profits are proportional (by a factor of $\gamma > 0$) in the two variations, it is intuitive that the set of equilibrium strategies will be the same. We also find an equivalence with our base model.

Proposition B.3 *The equilibrium of the biform and the two-stage variations are unique, and the set of equilibrium strategies satisfies* $\mathcal{V}^{B} = \mathcal{V}^{TS} = \mathcal{V}$.

Proof Given the equilibrium definitions, we have that $\mathcal{V}^{\mathrm{TS}}$ satisfies the following two conditions: if $\mathbf{s}_i \in \mathcal{V}^{\mathrm{TS}}$ then $\Pi_i^{\mathrm{TS}}(\mathcal{V}^{\mathrm{TS}}) > 0$; and if $\mathbf{s}_i \notin \mathcal{V}^{\mathrm{TS}}$ then $\Pi_i^{\mathrm{TS}}(\mathcal{V}^{\mathrm{TS}} \cup \mathbf{s}_i) \leq 0$. Similarly, if $\mathcal{V}^{\mathrm{B}}$ satisfies $\mathbf{s}_i \in \mathcal{V}^{\mathrm{B}}$ then $\Pi_i^{\mathrm{B}}(\mathcal{V}^{\mathrm{B}}) > 0$ and if $\mathbf{s}_i \notin \mathcal{V}^{\mathrm{B}}$ then $\Pi_i^{\mathrm{B}}(\mathcal{V}^{\mathrm{B}}) \leq 0$. The proof has three steps.

[14] Although the equilibrium profits are unique, there are multiple possible equilibrium price vectors. As long as the firm with the second highest WTP prices at marginal cost, the other losing firms can price higher.

(i) First we establish equivalence between the biform and the two-stage variations ($\mathcal{V}^{\mathrm{B}} = \mathcal{V}^{\mathrm{TS}}$). Since $\Pi_i^{\mathrm{B}}(\mathcal{E}) = \gamma \Pi_i^{\mathrm{TS}}(\mathcal{E})$, it follows that $\mathcal{V}^{\mathrm{B}} = \mathcal{V}^{\mathrm{TS}}$. In other words, the two equilibrium strategy sets are equivalent.

(ii) Next we show that any viable strategy in $\mathcal{V}$ is an equilibrium strategy in the two variations ($\mathcal{V} \subseteq \mathcal{V}^{\mathrm{TS}}$). Suppose that $\mathbf{s}_i \in \mathcal{V}$. Then $AV_i(\mathcal{S}) > 0$ by definition of $\mathcal{V}$. Because $AV_i(\mathcal{E}) \geq AV_i(\mathcal{S})$ for any $\mathcal{E} \subseteq \mathcal{S}$, we have that $AV_i(\mathcal{V}^{\mathrm{TS}} \cup \mathbf{s}_i) > 0$ and hence $\mathbf{s}_i \in \mathcal{V}^{\mathrm{TS}}$.

(iii) Finally, we show that all equilibrium strategies from the variations are viable ($\mathcal{V}^{\mathrm{TS}} \subseteq \mathcal{V}$). The proof is by contradiction. Suppose that $\mathbf{s}_i \in \mathcal{V}^{\mathrm{TS}}$ and $\mathbf{s}_i \notin \mathcal{V}$. Then $AV_i(\mathcal{V}^{\mathrm{TS}}) > 0$ and $AV_i(\mathcal{S}) = 0$. For both conditions to hold simultaneously, there must be at least one strategy in $\mathcal{S} \notin \mathcal{V}^{\mathrm{TS}}$ that keeps $\mathbf{s}_i$ from having added value. Thus, there is a customer α and a strategy $\mathbf{s}_j \notin V^{TS}$ for whom $\mathbf{s}_i$ has the maximum value creation for all $\mathbf{s} \in \mathcal{V}^{TS}$ but $\mathbf{s}_j$ has greater value creation.[15] But this implies that $AV_j(\mathcal{V}^{\mathrm{TS}} \cup \mathbf{s}_j) > 0$, which contradicts $\mathbf{s}_j \notin \mathcal{V}^{\mathrm{TS}}$.

We conclude that $\mathcal{V}^{\mathrm{TS}} = \mathcal{V}$. ∎

We have thus demonstrated that the simple model described in the paper's main text—with entry into all positions and market shares determined by value creation for each customer—is equivalent to two rational entry models with low barriers to entry and high levels of rivalry.

[15]We do not consider the possibility that they have the same value creation (i.e., $v(\mathbf{s}_j, \alpha) = v(\mathbf{s}_i, \alpha)$) because this happens with probability zero given that the contribution functions underlying the attribute levels are drawn from continuous distributions.

References

Barnett, V., T. Lewis. 1994. *Outliers in Statistical Data*. 3rd ed. Wiley, New York, NY.

Baumann, O., N. Siggelkow. 2010. Complexity and competitive advantage. P. Allen, S. Maguire, B. McKelvey, eds., *SAGE Handbook of Complexity and Management*. SAGE Publications, Thousand Oaks, CA. Forthcoming.

Ben Akiva, M., A. De Palma, J. F. Thisse. 1989. Spatial competition with differentiated products. *Regional Science and Urban Economics* **19**(1) 5–19.

Brandenburger, A. M., H. Stuart. 1996. Value-based business strategy. *Journal of Economics & Management Strategy* **5**(1) 5–24.

Brandenburger, A. M., H. Stuart. 2007. Biform games. *Management Science* **53**(4) 537–549.

Canoy, M., M. Peitz. 1997. The differentiation triangle. *Journal of Industrial Economics* **45**(3) 305–328.

Casadesus-Masanell, R., J. E. Ricart. 2010. From strategy to business models and onto tactics. *Long Range Planning* **43**(2-3) 195–215. Special Issue.

Chatain, O., P. Zemsky. 2007. The horizontal scope of the firm: Organizational tradeoffs vs. buyer-supplier relationships. *Management Science* **53**(4) 550–565.

Csaszar, F. A., N. Siggelkow. 2010. How much to copy? Determinants of effective imitation breadth. *Organization Science* **21**(3) 661–676.

De Palma, A., V. Ginsburgh, Y. Y. Papageorgiou, J. F. Thisse. 1985. The principle of minimum differentiation holds under sufficient heterogeneity. *Econometrica* **53**(4) 767–781.

De Palma, A., V. Ginsburgh, J. F. Thisse. 1987. On existence of location equilibria in the 3-firm Hotelling problem. *Journal of Industrial Economics* **36**(2) 245–252.

Durrett, L., V. Limic. 2003. Rigorous results for the NK model. *Annals of Probability* **31**(4) 1713–1753.

Gavetti, G., D. A. Levinthal, J. W. Rivkin. 2005. Strategy making in novel and complex worlds: The power of analogy. *Strategic Management Journal* **26**(8) 691–712.

Hotelling, H. 1929. Stability in competition. *The Economic Journal* **39**(153) 41–57.

Kauffman, S. A. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, New York, NY.

Lancaster, K. J. 1966. New approach to consumer theory. *Journal of Political Economy* **74**(2) 132–157.

Lancaster, K. J. 1990. The economics of product variety: A survey. *Marketing Science* **9**(3) 189–206.

Lenox, M. J., S. F. Rockart, A. Y. Lewin. 2006. Interdependency, competition, and the distribution of firm and industry profits. *Management Science* **52**(5) 757–772.

Lenox, M. J., S. F. Rockart, A. Y. Lewin. 2007. Interdependency, competition, and industry dynamics. *Management Science* **53**(4) 599–615.

Levinthal, D. A. 1997. Adaptation on rugged landscapes. *Management Science* **43**(7) 934–950.

MacDonald, G., M. D. Ryall. 2004. How do value creation and competition determine whether a firm appropriates value? *Management Science* **50**(10) 1319–1333.

Motta, M. 1993. Endogenous quality choice: Price versus quantity competition. *Journal of Industrial Economics* **41**(2) 113–131.

Nelson, R. R. 1991. Why do firms differ, and how does it matter? *Strategic Management Journal* **12**(Special Issue: Fundamental Research Issues in Strategy and Economics) 61–74.

Porter, M. E. 1985. *Competitive Advantage: Creating and Sustaining Superior Performance*. Free Press, New York, NY.

Porter, M. E. 1996. What is strategy? *Harvard Business Review* **74**(6) 61–78.

Rhee, B. D. 1996. Consumer heterogeneity and strategic quality decisions. *Management Science* **42**(2) 157–172.

Rivkin, J. W. 2000. Imitation of complex strategies. *Management Science* **46** 824–844.

Rivkin, J. W., N. Siggelkow. 2003. Balancing search and stability: Interdependencies among elements of organizational design. *Management Science* **49**(3) 290–311.

Rivkin, J. W., N. Siggelkow. 2007. Patterned interactions in complex systems: Implications for exploration. *Management Science* **53**(7) 1068–1085.

Saloner, G., A. Shepard, J. M. Podolny. 2001. *Strategic Management*. John Wiley, New York, NY.

Shaked, A., J. Sutton. 1982. Relaxing price competition through product differentiation. *Review of Economic Studies* **49** 3–13.

Siggelkow, N., D. A. Levinthal. 2003. Temporarily divide to conquer: Centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation. *Organization Science* **14**(6) 650–669.

Sommer, S. C., C. H. Loch. 2004. Selectionism and learning in projects with complexity and unforeseeable uncertainty. *Management Science* **50**(10) 1334–1347.

Spulber, D. F. 2004. *Management Strategy*. McGraw-Hill Irwin, Boston, MA.

Tirole, J. 1988. *The Theory of Industrial Organization*. MIT Press, Cambridge, MA.

Vives, X. 1999. *Oligopoly Pricing: Old Ideas and New Tools*. MIT Press, Cambridge, MA.

Zott, C., R. Amit. 2008. The fit between product market strategy and business model: Implications for firm performance. *Strategic Management Journal* **29**(1) 1–26.