12-2011

From primordial soup to platform-based competition: Exploring the emergence of products, systems, and platforms

C. Jason WOODARD
Singapore Management University, jason.woodard@olin.edu

Eric K. CLEMONS
University of Pennsylvania

Proceedings of the 44th Hawaii International Conference on System Sciences - 2011
http://doi.org/10.1109/HICSS.2011.208

From Primordial Soup to Platform-Based Competition: Exploring the Emergence of Products, Systems, and Platforms[i]

C. Jason Woodard jwoodard@smu.edu.sg
School of Information Systems, Singapore Management University

Eric K. Clemons clemons@wharton.upenn.edu
The Wharton School, University of Pennsylvania

Abstract

We use an agent-based NK model to explore the conditions under which standard platforms emerge among competing products. Our findings were inconclusive. We find that the usual Darwinian conditions needed for the emergence of complexity are sufficient to yield a limited reliance upon platforms with a core of common components, simply because evolution causes the population to converge on a set of products that contain combinations that "work well," yielding what we call "coincidental platform emergence." Economies of scale yield more use of common components, or "production platform emergence." *Positive participation externalities initially induce the highest degree of platform emergence through "usage platform emergence,*" *but this rapidly degenerates into simple monoculture. We find that lock-in, or freezing on early designs, can occur when variants arrive dynamically and not all choices are initially available, but that the cost is always a small fraction of participation benefits. Finally, we provide some extensions to the NK framework that improve its ability to address issues in system design and the assembly of components into products.*

1. Introduction

1.1. Focus of the paper

Our primary intent with this paper is to explore the minimum set of conditions needed to explain the emergence of technological systems and platforms. We provide more precise definitions later, but we consider a ***system*** to be a product with value contributed in part by the interaction of its parts; that is, a system is not just a collection of components, but rather is an assembled product with additional value created by its assembly, so that the value of the product exceeds the sum of the values of its components. We consider a ***platform*** to have emerged when several products share the use of one or more common parts, and where that part or collection of parts has largely supplanted the use of alternatives that could have been chosen instead. We will refer to an individual part of a system as a *component*; alternative forms of the components will be called *variants*.

Our secondary intent was to explore the extensions needed to adapt the NK modeling framework to address the evolution of artifacts with a complex internal structure. The NK framework was designed to explore the search for high-value "peaks" on a rugged "fitness landscape" whose dimensions correspond to interdependent design choices. These design choices are typically modeled as "bit strings" in which a complete design is fully described by a sequence of binary values (0 or 1). These values are normally treated symmetrically; there is no special significance to a 0 or a 1. To model the emergence of complex products, we imposed a more specific interpretation: 0 represents the absence of a given component, while 1 represents its presence (and higher values represent the presence of alternative variants). This additional structure requires some extensions to the basic NK model. We attempted to minimize these extensions, as explained below. Our concluding section comments on their limitations and additional extensions that may be fruitful to explore.

In support of our primary objectives, we looked for minimal sets of conditions that yield the emergence of systems and platforms in a population of multi-component products. We consider a set of conditions to be *minimal* if it is (1) *sufficient* to produce the desired outcome with high probability, and (2) *necessary* in that removing any of the conditions results in the desired outcome no longer being observed with high probability.

These simple investigations have a wide range of possible applications in the study of technology and industry evolution: (1) What are the conditions needed for complex products and systems to emerge? (2) When does a dominant set of core components emerge as a platform, displacing substitute components and supporting a diverse ecosystem of complementary ones? (3) Is the emergence of platforms predictable or even inevitable under certain conditions? (4) When a platform does emerge, is there a significant element of uncertainty or path dependence in the process by which it is selected? (5) To what extent does this process yield technological lock-in, and what are its costs? These questions are increasingly drawing the attention of scholars at the intersection of product design [2, 17, 21], industrial economics [3, 19, 22] and competitive strategy [7, 9, 11].

1.2. Modeling technique and objectives

Our modeling technique is based on the NK framework developed by Kauffman [14], who first used it to study biological evolution. He and others have successfully applied NK modeling to study organizational evolution [16, 20, 23] and technological innovation [8, 10, 14].

Our simulations assume a population of products assembled from a “primordial soup” of primitive components. The value (or fitness) of a product depends on the particular combination of components it contains. Boundedly rational product designers modify these combinations in search of higher fitness. In most of our experiments, a population-level selection process weeds out designers with inferior products and replaces them with new entrants. Our model thus includes both blind (or myopic) variation and selective retention, the key elements of evolution in both natural and artificial systems [4].

We began this investigation with six hypotheses on the emergence of systems and platforms.

The systems hypotheses:

Hypothesis 1 {Emergence of systems}: The evolutionary process simulated by the model will yield complex multi-component systems, i.e., products that are not merely collections of components but offer superadditive value.

Hypothesis 2 {Early adoption of core components}: The systems that arise will initially exhibit high pleiotropy [1, 18], which will quickly rise, and then decrease gradually over time. Systems will quickly adopt a few core components that affect many functions, and then add refinements and embellishments to successful designs later. Just as genes with high pleiotropy code for proteins that have widespread biological implications throughout the organism, components with high pleiotropy interact with many other components. Adding one of these “core” components can quickly lead to great increases in fitness, hence their frequent early inclusion in designs that subsequently prove successful. Adding a high-pleiotropy component to a mature system, on the other hand, tends to produce conflicts with existing components [10]. For this reason, we expect that components added later will tend to be lower in pleiotropy, yielding a decline after the initial peak.

The platform hypotheses:

Hypothesis 3 {Non-emergence of platforms in the basic model}: The evolutionary process examined in the first two hypotheses will not give rise to platforms with high probability.

Hypothesis 4 {Sufficient conditions for the emergence of platforms}: Platforms will emerge with high probability, but only after imposing conditions that reward designers for focusing on specific component variants in preference to others. The two most promising conditions that could lead to designers favoring certain components or sets of components are ***economies of scale*** and ***positive participation externalities***, also called network effects or participation benefits. When the use of specific variants by several designers results in lower costs for those variants, we will expect them to gain in popularity, even when there may have been no reason other than chance for their initial widespread adoption. When the use of the combination of several specific variants results in their widespread adoption, this may in some instances, for some forms of products, result in significant benefits, quite separate from the lower costs associated with the widespread adoption of their individual components. These benefits generally arise from interoperability of systems, ease of sharing work products, ease of obtaining trained and experienced workers, and, occasionally, from scope-based economies of scale that arise from sharing some additional products across a wider family of related platform members. We introduce both economies of scale and participation externalities with scalable parameters, so that we can investigate the extent to which they influence the emergence of platforms.

Hypothesis 5 {Role of chance in the emergence of platforms}: There is a considerable element of chance involved in selecting the variant(s) that emerge as the basis of a platform, just as there is chance in the evolution of biological forms [12].

Hypothesis 6 {Possibility of lock-in in the emergence of platforms}: There is a possibility that a dominant, stable platform emerges even though it is not the best that could have been obtained under the same environmental conditions.

1.3. Structure of the paper

The structure of the paper is as follows. Section 2 provides a brief review of the literature on components, products, systems, and platforms. Section 3 briefly reviews the literature on NK simulation, both in the context of its original use in modeling biological evolution, and in the context of its more recent use to model organizational evolution and learning. Section 4 explains the areas in which we extended the basic NK model to address the emergence of products, systems, and platforms. Section 5 restates our hypotheses more formally, and reviews the simulation methodology we followed to investigate them. Section 6 explores our findings and the extent to which our experimental results supported our hypotheses. Section 7 reviews our principal conclusions, explores the limitations of the current research and ways to address them, and outlines future research opportunities.

2. Literature review: components, products, systems, and platforms

Following the product design literature [5, 24], we view a ***product*** as a collection of ***components***. Not all such assemblies are "good," and not all good assemblies are "equally good." For a non-trivial product to exist, one that incorporates more than one component, the assembly has to possess more value than any of its individual parts. In a market-based economy, where individual components have either explicit prices or opportunity costs for their use, we would expect a component to be incorporated into an assembled product only if its contribution exceeds either its price or the opportunity costs associated with its use elsewhere. In the "primordial soup" of early biological evolution, on the other hand, it is customary to assume that the organic chemicals that were the precursors for life existed without scarcity.

Following the biological tradition, we focus on the value created by combining an unlimited supply of components into a finite population of products. In particular, we look for well-architected ***systems*** in which the components work together to create more value than they could create separately. Systems are characterized by superadditive value functions; the value of a system exceeds the sum of the values of its individual components (when these components are valued as simple single-component products).

Most modern designers start with one or more core components when designing a system [5]. The engine and drive train of a car, or the processor and bus when designing a laptop computer are core components. These components have a high degree of interaction with other components, both facilitating incremental design and constraining future changes [18]. Non-core components, in contrast, may still contribute considerable value, but can be added, removed, or changed without requiring major redesign of the entire product.

We view a product population as exhibiting the emergence of a ***platform*** if there is a common core component or set of components employed by a large fraction of the population. That is, platforms can be detected by the high common use of a small number of component variants.

What we call a variant is analogous to an allele in genetics. While the NK modeling literature typically assumes the existence of only two alleles per "gene" or design decision (noting that this is without loss of generality since more complex genes or decisions can be broken down and encoded as binary sequences), we explicitly allow multiple variants for a single type of component. Moreover, in some of our experiments we allow components to differ in the number of variants that are available. These variants arrive stochastically over time, enlarging the product design space in the same way that innovation enlarges the design space of products in the real world. As in the real world, variants are not necessarily functionally equivalent; some may offer higher performance or better compatibility with other components than others. However, we do not consider gradual evolution within a family of variants. The Intel 80386 interacted with other system components in much the same way as its predecessor the 80286, whereas in our model the characteristics of later arrivals are uncorrelated with their predecessors.

3. Literature review: the NK modeling approach to design evolution

3.1. The basic NK model

The NK family of models is highly stylized and does not support detailed descriptions of biological processes or organizational interactions. The limitations of the NK framework are also its strengths. It makes no attempt to embrace the domain-specific details of any specific setting; it does not attempt to predict the weather, the success of a product, or even the outcome of a generic marketing strategy. But because of this it can model a wide range of phenomena related to evolution, search, and learning. NK models have delivered qualitatively robust insights about basic evolutionary forces [1, 13], decision-making behavior in organizations [6, 20, 23], and technological innovation [10].

At the heart of an NK-style model is a fitness landscape that can be fully described by (1) N, the number of design or decision variables (yielding an N+1-dimensional landscape); (2) K, the complexity of interactions among the N variables, specifically the number of variables that jointly determine the "height" or fitness of a position on the surface of the landscape, and (3) P, a function that assigns a fitness value to each position (typically at random).

Each of the N variables represents a set of alternative decisions or design choices. In the typical case of binary variables, each position on the landscape corresponds to one of 2^N possible strategies or designs. A strategy or design is completely specified by its N-dimensional choice vector, and the fitness or payoff associated with that combination of choices is completely determined by its corresponding location on the fitness landscape.

Since high-dimensional hypercubes are difficult to visualize or to draw, it is customary to think of K as a measure of cragginess on a two-dimensional map of the N-dimensional space, with high K producing craggy and complex environments.

Evolution in NK models occurs at the level of the individual and sometimes also the population. At the individual level, agents (in our case representing product designers) repeatedly move to new positions on the landscape and are rewarded with higher payoffs if they move to higher-fitness positions, which represent strategies or designs that are better adapted to their environment. At the population level, we follow Levinthal [16] in modeling the entry and exit of designers according to their relative fitness; details are provided in section 5.1.

Since changes that decrease fitness tend to lead to the elimination of agents from the population, it is customary to assume that only beneficial changes are made. These changes occur via two mechanisms, short and long jumps. Short jumps involve changing a single decision or design variable, and thus moving to an adjacent position on the landscape. The direction of the change can be determined by chance or by various forms of directed (gradient) search; we explored a variety of decision rules, as described in section 5.1. As is well known, gradient search may stall on a local optimum, never reaching the global maximum. Many NK models also permit long jumps, which occur much less frequently in nature. In some models long jumps occur only after local optima are achieved, in others at random intervals, but these involve changing a large number of variables in an attempt to reach a remote local optimum that is higher than the current location.

3.2. Limitations of the basic NK model for modeling complex system designs

In the basic NK model there are N “bits” (design dimensions), each of which also represents a function that contributes to fitness. All functions are influenced by K bits, and all bits influence the same number of functions. We know this is not true in biological systems. Some genes participate in a large number of functions (high pleiotropy); for example, the genes that regulate the core processes of DNA and RNA transcription affect nearly all cell functions. Likewise, some functions or traits are influenced by a large number of genes (high polygeny); the height of an organism, for example, depends on a large number processes that are regulated by the interaction of a large number of genes. Similarly, some organizational decisions have wide-ranging implications (mimicking biological pleiotropy) and some organizational functions are dependent upon a large number of decisions (mimicking biological polygeny). The choice of a processor for a laptop will to some extent constrain choices of operating system, which in turn will constrain other choices in available software. The processing speed of a laptop will likewise be determined by choice of processor, number of processors, size and speed of memory, and perhaps size, speed, and number of disk drives.

These limitations were recognized by researchers in biology [1] and technological innovation [10], who responded by relaxing the requirement that K be constant across all elements in a genome or product design. Instead, they allowed elements to vary in their pleiotropy, making it possible to distinguish between core and peripheral components in an organism, organization, or product. We were not the first to take this important step, but we rely upon it in our extensions to the NK framework.

4. Extensions to the NK framework

To adapt the NK framework to study the emergence of complex designed artifacts, we needed to make three additional changes:

First, we focus on a particular kind of design decision, namely the inclusion or exclusion of a component in a product. Instead of N binary decisions, we consider products containing up to C types of components. A C-dimensional “bit string” thus represents a subset of components that are selected for a given product, or equivalently one of the 2^C possible products that can be assembled from a set of C components. Following Altenberg [1] and Frenken [10], we assume that each component affects k out of F functions, where F is a constant and k is a number between 1 and F that is drawn independently at random for each component. The functions affected by a given component are called its pleiotropy set, while the components that affect a given function are called its polygeny set. The fitness contribution of each function is determined by the presence or absence of the components in its polygeny set, As in the basic NK model, every combination of variants (i.e., every unique bit string) in a function’s polygeny set is associated with a different fitness contribution drawn independently from a uniform distribution on the unit interval. If no components in a function’s polygeny set are present, its fitness contribution is zero. The total fitness of a product is the average of the F fitness contributions.

Second, as indicated above, we allow the possibility of multiple variants of a given component type (e.g., Motorola and Intel processors). This is similar to the possibility of having more than two alleles for a given gene, which is also provided in the basic NK model. We go beyond the basic model in allowing the number of variants per component to differ by component. While most of our simulations assume a constant number of variants per component, V, this additional flexibility allows us to model the arrival of variants over time, so that some components may have more variants than others in a given period.

Third, we needed to make a slight change to the nature of selection in competition between designers. While we mainly adopt the assumptions made by Levinthal [16], we assume that new entrants arrive with “empty” products (with zero components and thus a value of zero) rather than products containing a random subset of components. In contrast to organizational evolution, where organizations like firms tend to enter the competitive marketplace with the same set of “components” (e.g., some form of accounting department, sales and distribution channels, and decision-making mechanisms) even if they organize these components in novel ways, product designers tend to either imitate an existing design or start with a clean slate and build incrementally. Seldom if ever does a designer create a product by picking up a full set of components chosen at random and trying to assemble them into a functional whole.

5. Experimental design

5.1. General review of methodology

To examine our hypotheses, we executed a series of computational experiments. For each combination of parameter settings, we ran 200 independent trials using different random seeds, allowing the effects of each combination to be observed with reasonably low variance. Each seed gave us a run with a unique randomly selected competitive landscape or environment, a randomly selected initial set of components and variants, and an initially empty set of competing designs. We allowed each experimental run to continue for 100 time periods. We initially calibrated our experiments using run lengths of 100, 200, 500, and 1,000, and found the results to be largely stable after 100 periods.

To explore the effects of search and selection, we simulated nine distinct evolutionary processes:

(1) Blind chance with selection. In each period the existing set of designers can make one change to their products, and all changes are made at random. A component can be added, or, for products that already have at least one component, a component can be dropped. Alternatively, any single component can be replaced with an alternative variant. Following Levinthal [16], designers survive with probability proportional to the value of their products relative to the highest-value product in the population. Non-survivors are replaced either with new entrants (whose products are initially empty) or entrants who replicate existing product designs, with the probability of new entry equal to the ***genetic load*** of the population [25]. These assumptions yield dynamics that resemble, in a stylized way, those of biological, Darwinian evolution.

(2) Myopic 1-step adaptation without selection. In this process, designers make the first change they find, selected at random, that leads to an improvement in their product’s value (fitness), making at most one change at a time. Natural selection does not cull out the least fit, but all designers do make directed progress toward local peaks on their fitness landscape. Consistent with all hill-climbing algorithms, there is the possibility that designers will reach and get “stuck” at a local peak that is not a global maximum; this is true in each of the eight variations on this process described below.

(3) Myopic 2-step adaptation without selection. This is the same as myopic 1-step without selection except that the designer considers making one or two changes in a single time period and stops when a change is found that leads to an increase in fitness.

(4-5) Greedy 1- and 2-step adaptation without selection. This is the same as processes 2 and 3 except that the designer examines all possible 1- and 2-step changes and makes the change that produces the largest increase in fitness. The process is repeated until no more fitness-increasing steps are found.

(6-9) Myopic 1- and 2-step adaptation with selection; greedy 1- and 2-step adaptation with selection. These are the same as processes 2–5 except that natural selection (as described in process 1) is implemented in addition to the indicated mode of hill-climbing search.

5.2. Implementation of hypotheses

Hypothesis 1 was the first of our systems hypotheses, suggesting that systems, or products with superadditive value, would emerge when designers pursued local improvement. Hypothesis 2 suggested that high-pleiotropy components would tend to be adopted first when systems emerged. Testing these hypotheses involved examining runs with all eight combinations of myopic and greedy design, one- and two-step walks, with and without selection. We checked to see when superadditive value emerged, to determine if the emergence of superadditive value entailed the early adoption of high-pleiotropy components, and to determine which conditions were necessary for the emergence of superadditive value.

Hypothesis 3 and 4 were the first of our platform hypotheses. Hypothesis 3 suggested that platforms, unique combinations of specific variants of components that were widely adopted, would not emerge under the conditions we used to examine hypotheses 1–2; testing simply involved checking for the presence of such combinations. More specifically, our definition of a platform required examining not only products but also the distribution of products in their environment. A platform was said to emerge when (1) there was a core combination of components that existed with common choice of variants, and (2) there was a significant number of products that used the same combination of core products and core variants of these products, (3) there was a diversity of choices for other components employed in the set of products using the common, and (4) the core combination was high pleiotropy, and not merely a combination of components that co-occurred by chance. Conditions (1) and (2) require that there be a combination of choices that is common to several products, such as the choice of an Intel chip and a Microsoft operating system for a platform to have emerged. Condition (3) demands that declaring a platform to have emerged requires that the core combination support a diverse combination of products, rather than declaring the emergence of a platform when a single design outcompetes all others. Condition (4) ensures that the core of a platform exhibits the characteristics that we intuitively expect to see in a platform, that is, a core that is highly related and tightly coupled.

Testing hypothesis 3 involved examining the set of products that emerged at the end of our runs for test of hypothesis 2, to see if platforms had or had not emerged. Hypothesis 4 suggested that platforms would emerge when economies of scale or positive participation externalities were added, and testing involved repeating the runs used to test hypothesis 2, but with the addition of various levels of economies of scale or participation externalities. Hypotheses 5 and 6 addressed the possibility of path dependence and lock-in, and required allowing variants to arrive dynamically over time. We tested the values of the products that emerged at the end of runs using both static and dynamic variant arrival, where static arrival had all variants available at the start of the run and dynamic arrival had some variants available at the beginning and had others arrive over time. When the values of the set of products using the core platform were largely the same with both static and dynamic arrival we concluded that no significant path dependency arose, testing hypothesis 5. When the values were significantly lower using dynamic arrival we concluded that lock-in had occurred, resulting in the widespread adoption of a platform that was available early but that was inferior to one that would become available only later with the arrival of one or more new variants, thus testing hypothesis 6.

6. Principal findings

6.1. Preliminary observations

We expected to find that complex products emerge under all of our simulated evolutionary processes, since they all include the basic elements of blind variation and selective retention that characterize evolution in biological, organizational, and technological settings. As expected, and consistent with the NK modeling literature, these processes proved sufficient to produce the emergence of complex product designs. We find that complexity emerges whenever combining components is permitted.

We ran these simulations under a range of conditions, including (1) myopic one-step search, (2) myopic two-step search, (3) greedy one-step search, and (4) greedy two-step search, with and without selection, while varying the number of components. For all runs the number of functions, components and variants was fixed at 16. As seen in figure 2, all results were similar except for blind search in a fixed population. Since the results are so similar, for the rest of the paper we use only one-step greedy search with natural selection. However, we noted that even in the absence of natural selection some increase in product complexity occurs, largely because there is no cost associated with addition of a component; this should have been an early indicator to us of problems we would encounter in our study of system and platform emergence. We also show that product complexity grows sharply as the number of components increases (allowing for greater complexity), but more slowly as the number of variants per component increases (allowing for better choices).

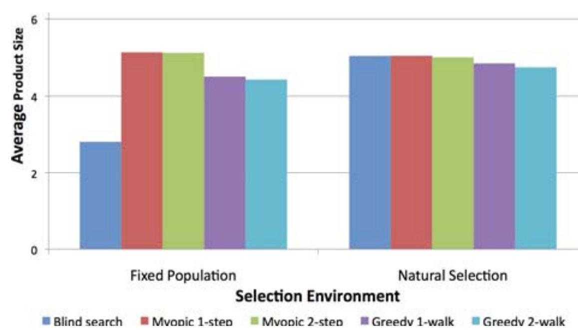

Figure 1. Average product size as a function of the number of components, with combinations of search mode and selection environment.

6.2. Examination of system hypotheses

We defined a system as a product whose components interact to produce superadditive value. Hypothesis 1 suggested that simple evolutionary dynamics would yield products with this property. The data do not support this hypothesis, as shown in figure 2. The figure actually uses a weaker concept of superadditivity than defined above; here a product is deemed to be superadditive if its value is *at least* that of the sum of its components. Even under this less stringent test (which single-component products pass by definition, yielding the ridge along the top left of the figure), superadditivity declines rapidly for higher numbers of components and essentially vanishes when the number of components available exceeds 3. In settings where 16 components are available, fewer than 1% of products in an average run are even weakly superadditive.

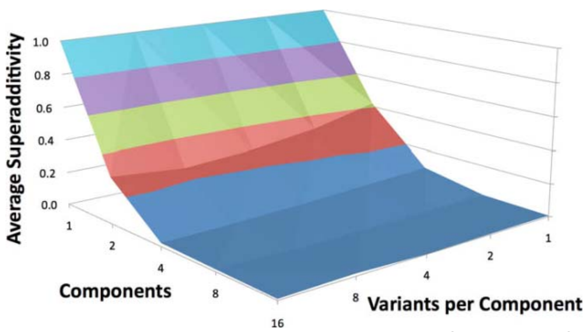

Figure 2. Average superadditivity as a function of number of functions and number of components.

Since systems did not emerge, hypothesis 2 is untestable. However, we do find that average pleiotropy in the general product population rises and then falls over time, consistent with prior results [1, 10].

Analysis makes it clear why the systems hypotheses were unlikely to be true. Since the basic NK model does not include the concept of a component or of complexity, it does not impose a cost on complexity or an opportunity cost on the inclusion of a component (and we did not introduce such a cost as an additional assumption). Our hypothetical example for this in our discussions was the concept of an asparagus-hat; if there is no cost to inserting a stalk of asparagus into the band of a hat, and if this increases the value of a hat even slightly, then under our model it will be done, even if the cost of the asparagus is less than the resulting increase in the value of the hat.

6.3. Examination of platform hypotheses

Our definition of a platform required examining not only individual products but also the distribution of products in their environment, requiring both widespread adoption of a common set of high-pleiotropy components and their use in a diversity of products. Contrary to our expectations, we found that a moderate degree of platform emergence was observed even under the conditions of hypothesis 1, as shown by the lines marked CRS in figure 3. This does not support hypothesis 3. Good fitness is rewarded with reproductive success, and populations tend to cluster around a limited number of local maxima; since products close to each other in a cluster will certainly share two or more components, they will appear to be based on a common platform, a phenomenon we will call "coincidental platform emergence." We next tested economies of scale and participation externalities to see if more extreme platform behavior was induced, and the results we obtained were also not entirely what we anticipated.

We added economies of scale and represented them by a reduction in cost for a component over time. We assumed that the cost savings from using the specific variant of any component was a function of the number that had been previously produced, with savings given by αe^{-cn}, where α and c represent scale factors, n represents the total number of the specific variant already produced. The more produced, the greater the benefit from using the specific variant. Economies of scale are represented as a decrease in cost of the specific component, leading to a corresponding increase in the total value of any product using that component. The presence of sufficiently large economies of scale does induce the emergence of platforms, as shown by the lines marked EOS in figure 3. This is the form of platform behavior exhibited by the wide range of GM cars that use a common engine or the number of aircraft built on the 747 airframe, and we therefore term it "production platform emergence." This is consistent with the economies of scale portion of hypothesis 4.

We also added positive network externalities (participation benefits). We assumed these were of the form $\beta\ n\ \log(n)$, where β is once again a scale factor and in this instance n is the total number of products in use that were built using the platform's combination of unique variants of specific components, not the total number of products that had ever been made that shared those components. It is customary to let participation benefits increase with the number of users n, and indeed to increase more rapidly than n; letting benefits vary with n^2 ran the risks of over-stating benefits and of rapidly swamping the other elements of the value function. We added these benefits to all pair-wise combinations in use. The presence of participation externalities is sufficient to induce the very rapid emergence of platforms, as shown by the early behavior of the lines marked PEP in figure 3. This strongly supports the participation externalities portion of hypothesis 4. Since this is based on the benefits that products contribute when different users can interact, we term this "usage platform emergence." However, we note that over time almost all variation disappears and the population converges to monoculture. This was as if ***all*** Apple products used the same chip and the same release of OS X, and all Windows machines likewise used the same chip and the same release of Windows; since there is no longer a diversity of products built around a common core one dimension of our platform measure, and hence the overall platform index score, ultimately collapses. This convergence is supported by examination of figure 4. Without variation in consumer preferences for tradeoffs between price and performance or power and portability, and without families of compatible variants, the rewards for ***exact*** replication are great enough to drive the products to uniformity.

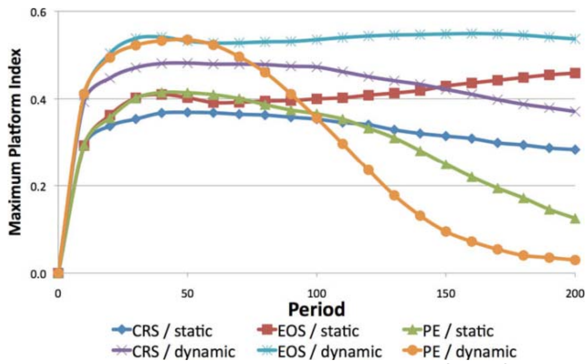

Figure 3. Examining the emergence of platforms when variants arrive over time.

Figure 3 shows how platforms emerge, with 16 functions, 16 components, selection, and greedy search, with and without economies of scale or participation externalities, and with and without arrival of variants over time. When allowing variants to arrive over time we examined both ***undirected*** innovation (in which all components were equally likely to experience innovative arrival of new variants), and ***directed*** innovation (in which the most widely used components are the ones that experience the arrival of the most new variants). Directed innovation increases the likelihood the more widely used components will see the arrival of more potential replacements, and hence even with randomly distributed attributes, more potential improvements. Figure 3 shows only the results with undirected innovation; these results are robust under directed innovation and changes to the number of variants.

We varied the parameters α and β, and as expected, these greatly influenced the results. For very low values of α and β platforms did not emerge, just as they did not emerge when these values were implicitly set to zero by omitting economies of scale and participation externalities. For very high values of α and β the environment very quickly selected for monocultures around core combinations of components instead of diversity.

Figure 5 shows the average product base value (net of scale or participation bonuses) with fixed and dynamic variable arrival. When variants arrive over time we found that the set of products that emerged did, as expected, have lower base values. There are differences, supporting hypothesis 5, and these differences produce costs, supporting hypothesis 6. This is true even without economies of scale or participation externalities; this is not surprising because some products may emerge as extremely fit using high pleiotropy components that emerged early and locked in on their use; when better alternatives arrived later it was not possible to incorporate them in place of early arrivals without simultaneously making too many other changes. With economies of scale the benefits of early arrival become larger, making it even more difficult for later variants to be considered and magnifying the effect. The effect is most pronounced in the presence of strong participation externalities, resulting in lock-in due to the success of an early variant that works, precludes adoption of later arrivals due to its own participation externalities, and results in stable but socially less attractive outcomes.

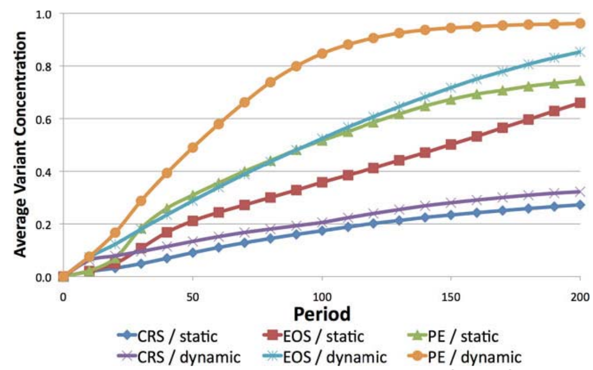

Figure 4. Examining the emergence of platforms when variants arrive over time.

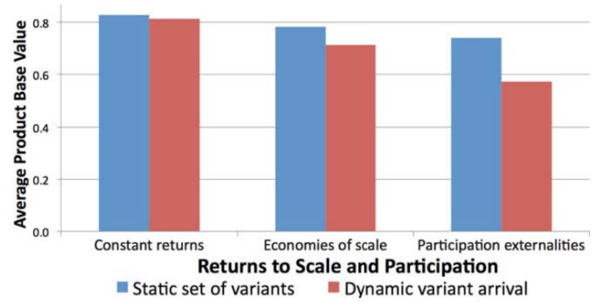

Figure 5. Examining the possibility of path dependence by comparing platform values when variants arrive over time.

Figure 6 examines the significance of the cost of lock-in. We conclude that a high cost of lock-in only occurs when the benefits from participation externalities are extraordinarily high. That is, although the costs of lock-in can be large (greater than 15% of the base value created by the system without participation externalities) this only occurs when the benefits from participation are extraordinarily high, dwarfing the value generated by the product itself (as highlighted by the marked rectangle in the upper right corner). This would self-evidently be true for email, for example, where the value of sending and receiving messages is far greater than the value of simply filing and organizing them. Even in this case, however, regulatory concern might not be warranted; in most cases, new products emerge that either inherent the participation externalities of previous platforms (vendors' new operating system releases usually are backwards compatible variants of a previous release) or are engineered to share those of a major competitor (Microsoft Word initially could read files produced by WordPerfect).

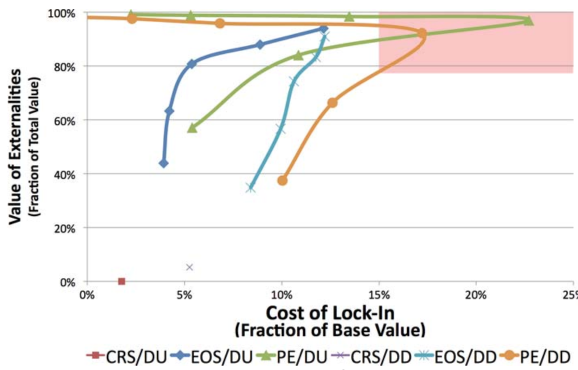

Figure 6. Examining cost of path dependency when variants arrive over time.

7. Conclusions

7.1. Summary of major findings

The fact that complexity arises even in the absence of natural selection is disturbing. The fact that complex products emerge under almost any conditions was our first indication that NK modeling does not inherently accommodate the concept of components, or of individual families of components and their variants, and that it may be an inappropriate tool for addressing some problems in design evolution. With no opportunity cost associated with combining products and no minimum activation energy or hurdle that must be overcome to achieve combination, complex products emerge under any and all conditions.

Our experience with hypotheses 1 and 2 may likewise reflect limitations in the traditional NK model when applied to this problem setting. As noted, these hypotheses were not supported, because we did not observe the emergence of systems with super-additive value. This runs counter to our intuition. An organism is more fit than the sum of its individual organs or organelles, and a product is more valuable than a stack of its individual components. With no opportunity cost associated with combining products and no hurdle that must be overcome to achieve combination, complex products do indeed emerge under any and all conditions, and they continue adding components until the marginal value of the last addition drops to zero and until the superadditivity of the assembly is eliminated.

Contrary to our expectations (and to hypothesis 3), we observed moderate platform emergence even under the conditions of the basic model. This effect was strengthened by the addition of economies of scale and participation externalities, weakly supporting hypothesis 4. Consistent with hypotheses 5 and 6, we found that some path dependence exists, but the cost of lock-in is either small or swamped by the benefits of participation externalities.

7.2. Limitations of present research

Our definition of a component variant is too limited, and will be extended to include families of variants, where all members of a family can contribute equally in participation externalities.

Likewise, our definitions of competitive landscape and fitness are too limited. We show all products competing on the same landscape. In fact there are niches, if for no other reason than not all consumers have the same preferences for product attributes, or for product attributes and price.

7.3. Extensions and future research

We plan to extend the concept of a variant to include component families, with a broader definition of participation externalities. We will also add niches based on heterogeneous consumer behavior, allowing products that occupy sufficiently distant places in the landscape to do so without intense competition, reflecting differences in consumers' preferences. We would also like to model the balance of cooperation and competition between platform vendors and the smaller firms in their business ecosystems.

More generally, we see opportunities to study a wide range of issues relating to product architecture and competitive strategy that are intractable using conventional analytic modeling tools, especially those involving dynamic behavior over time. These issues are especially salient in platform industries, where firms grapple with complex decisions that span the domains of system design and economic incentives. This work has suggested a small set of augmentations to the Altenberg–Frenken branch of the NK family that will facilitate the application of NK modeling to the study of product and system design:

1. The concept of ***component***, and component ***variants***, as modeled in this work.
2. The concept of variant ***families***, as proposed in our extensions.
3. The concept of ***component opportunity cost***, to ensure the absence of mindless accretion.
4. The concept of platform benefits, arising from the reduced opportunity cost as component usage increases, and the positive participation externalities that may arise from ecosystems built around common components or sets of components.

8. References

[1] L. Altenberg, "Evolving Better Representations Through Selective Genome Growth," in *Pro-*

ceedings of the IEEE World Congress on Computational Intelligence, Part 1. IEEE, Piscataway NJ, pp. 182-187, 1994.

[2] C. Y. Baldwin and K. B. Clark, *Design Rules, Volume 1: The Power of Modularity*, MIT Press, Cambridge MA, 2000.

[3] T. F. Bresnahan and S. Greenstein, "Technological Competition and the Structure of the Computer Industry," *Journal of Industrial Economics*, vol. 47, pp. 1-40, 1999.

[4] D. T. Campbell, "Variation and Selective Retention in Socio-Cultural Evolution," in H. R. Barringer et al. (eds), *Social Change in Developing Areas: A Reinterpretation of Evolutionary Theory*. Schenkman, Cambridge MA, 1965.

[5] K. B. Clark, "The Interaction of Design Hierarchies and Market Concepts in Technological Evolution," *Research Policy*, vol. 14, pp. 235-251, 1985.

[6] F. Csaszar and E. Clemons, "Governance of the IT Function: Valuing Agility and Quality of Training, Cooperation and Communications," in *Proceedings of the 39th Hawaii International Conference on System Sciences* (HICSS'06), IEEE, Piscataway NJ, 2006.

[7] T. R. Eisenmann, "Managing Shared and Proprietary Platforms," *California Management Review*, vol. 50, no. 4, pp. 31-53, 2008.

[8] S. K. Ethiraj and D. Levinthal, "Modularity and Innovation in Complex Systems," *Management Science*, vol. 50, pp. 159-173, 2004.

[9] D. S. Evans, A. Hagiu and R. Schmalensee, *Invisible Engines: How Software Platforms Drive Innovation and Transform Industries*, MIT Press, Cambridge MA, 2006.

[10] K. Frenken, *Innovation, Evolution and Complexity Theory*, Edward Elgar, Cheltenham UK, 2006.

[11] A. Gawer and M. A. Cusumano, *Platform Leadership: How Intel, Microsoft, and Cisco Drive Industry Innovation*, Harvard Business School Press, Boston, 2002.

[12] S. J. Gould, *Wonderful Life: The Burgess Shale and the Nature of History*, W. W. Norton, New York, 1989.

[13] S. Kauffman and S. Levin, "Towards a General Theory of Adaptive Walks on Rugged Landscapes," *Journal of Theoretical Biology*, vol. 128, pp. 11-45, 1987.

[14] S. Kauffman, J. Lobo and W. G. Macready, "Optimal Search on a Technology Landscape," *Journal of Economic Behavior and Organization*, vol. 43, pp. 141-166, 2000.

[15] M. W. Kirschner and J. C. Gerhart, *The Plausibility of Life*, Yale University Press, New Haven, 2005.

[16] D. A. Levinthal, "Adaptation on Rugged Landscapes," *Management Science*, vol. 43, pp. 934-950, 1997.

[17] M. H. Meyer and A. P. Lehnerd, *The Power of Product Platforms: Building Value and Cost Leadership*, Free Press, New York, 1997.

[18] J. P. Murmann and K. Frenken, "Toward a Systematic Framework for Research on Dominant Designs, Technological Innovations, and Industrial Change," *Research Policy*, vol. 35, pp. 925-952, 2006.

[19] G. G. Parker and M. W. Van Alstyne, "Two-sided Network Effects: A Theory of Information Product Design," *Management Science*, vol. 51, pp. 1494-1504, 2005.

[20] J. W. Rivkin, "Imitation of Complex Strategies," *Management Science*, vol. 46, pp. 824-844, 2000.

[21] D. Robertson and K. Ulrich, "Planning for Product Platforms," *Sloan Management Review*, vol. 39, no. 4, pp. 19-31, 1998.

[22] J.-C. Rochet and J. Tirole, "Platform Competition in Two-sided Markets," *Journal of the European Economic Association*, vol. 1, pp. 990-1029, 2003.

[23] N. Siggelkow and J. W. Rivkin, "Speed and Search: Designing Organizations for Turbulence and Complexity," *Organization Science*, vol. 16, pp. 101-122, 2005.

[24] K. Ulrich, "The Role of Product Architecture in the Manufacturing Firm," *Research Policy*, vol. 24, pp. 419-440, 1995.

[25] E. Wilson and W. Bossert, *A Primer of Population Biology*, Sinauer Associates, Sunderland MA, 1971.

[i] Research supported in part by a grant from the Wharton-SMU Research Centre.