Strategy Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Searching for Great Strategies

Thomas Fink, Pankaj Ghemawat, Martin Reeves

To cite this article:
Thomas Fink, Pankaj Ghemawat, Martin Reeves (2017) Searching for Great Strategies. Strategy Science 2(4):272-281. https://doi.org/10.1287/stsc.2017.0052

informs
http://pubsonline.informs.org/journal/stsc/

STRATEGY SCIENCE
Vol. 2, No. 4, December 2017, pp. 272–281
ISSN 2333-2050 (print), ISSN 2333-2077 (online)

Searching for Great Strategies

Thomas Fink,[a, b] **Pankaj Ghemawat,**[c, d] **Martin Reeves**[e]

[a] London Institute for Mathematical Sciences, Mayfair, London W1K 2XF, United Kingdom; [b] Centre National de la Recherche Scientifique, 75016 Paris, France; [c] Stern School of Business, New York University, New York, New York 10012; [d] IESE Business School, 08034 Barcelona, Spain; [e] BCG Henderson Institute, Boston Consulting Group, New York, New York 10001

Contact: tmafink@gmail.com, (TF); pghemawat@iese.edu, http://orcid.org/0000-0001-8633-0243 (PG); reeves.martin@bcg.com, (MR)

Received: September 14, 2017
Revised: November 8, 2017
Accepted: November 10, 2017

https://doi.org/10.1287/stsc.2017.0052

Abstract. We focus on answering the question posed for this special issue by elaborating a specific perspective, involving information-enabled search, in which firms add capabilities (or components) that expand what they can accomplish in the product market arena, and the key strategic choices concern the kinds of capabilities that are added. We establish that measuring how a new candidate component interacts with the components we already have can be a reasonable proxy for how they will combine with new components which we don't yet have. This allows us to compare the performance of "impatient strategies" focused on the current usefulness of a new component and "patient strategies" focused on anticipated long-term usefulness. Their relative performance depends on how far the innovation process has progressed, and on the structure of the innovation space itself. In particular, a flattening in the increase of complexity implies an increase in the relative attractiveness of patient strategies over impatient ones, i.e., constitutes a signal to a switch strategies. It is therefore possible to construct information-based adaptive search strategies, which outperform either random strategies or fixed (patient or impatient) strategies for component selection. And there are broader implications for strategy as well.

Funding: T. Fink acknowledges support from the European Commission FP7 [Grant 611272] under the Growth and Innovation Policy-Modelling project (GROWTHCOM).

Keywords: search • combinatorial innovation • irreversibility • uncertainty • adaptation • dynamic strategies

1. Introduction

Great strategies should be both impactful and innovative. But where do such strategic innovations come from? We explore the mathematics and empirics of a component-based model of innovation, and we propose a new perspective that, while recognizing that the search for innovations is inherently unpredictable, shows that it has mathematically predictable features. We show how these can be exploited to construct advantaged strategies that are contingent on both the characteristics of the search space and its degree of maturity.

This paper is organized as follows. Section 2 reviews the relevant literature and provides a rationale, from a comparative perspective, for our specific angle of attack. Section 3 presents illustrative examples that help make our ideas more concrete. Section 4 presents the basic model and some analytical results and corroborates them with historical data. Section 5 provides further discussion of the results and extensions, and Section 6 concludes.

2. Links to Literature

The question of where great strategies comes from has many answers, and there are theories and anecdotal "origin stories" to support each of them. But there is a fundamental tension between answers that emphasize favorable outcomes under conditions of uncertainty and those that assert intentionality (see Figure 1). Consider below thumbnail sketches of six different types of origin stories: while they are neither mutually exclusive nor exhaustive, they help characterize the contributions of this paper and relate them to the prior literature.

Luck, at the top of the wheel of fortune, underpins the simplest type of origin story. How we assess the contribution of luck depends on how far back we are willing to unravel success stories (Barney 1986). But systematic empirical analysis indicates that while luck matters, it is not the only factor behind performance differences between firms (e.g., Caves et al. 1977). This paper shows that what may appear to be serendipity actually has an underlying mathematical structure that can be exploited.

Foresight rubs shoulders with luck-based explanations, especially when the kind of foresight invoked involves intuition. The distinction between the two is sharpened by thinking of foresight as statistically significant information about what will happen: superior insight into how the future will unfold or a more complete picture of the space of possibilities. It is obvious that foresight can be the basis for superior performance under conditions of uncertainty. What is more challenging, however, is identifying prescriptive steps for how to achieve it. This paper sheds light on where such

Figure 1. The Wheel of Fortune

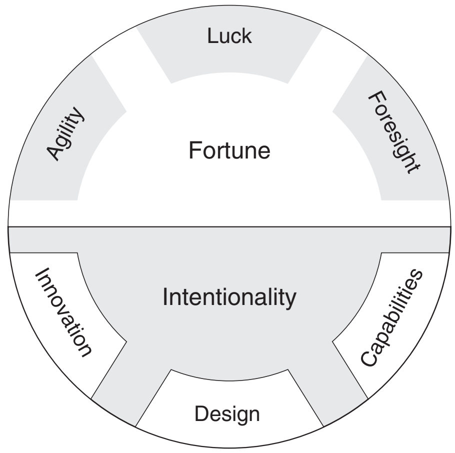

superior information might come from and how it can be exploited to formulate a farsighted strategy.

Capabilities-based thinking emphasizes heterogeneity in firms' opportunity sets. The analysis of capabilities has historically focused on the deepening of existing capabilities within predefined, well-bounded domains. But, as Pisano (2017) has recently stressed, there is an important distinction between deepening existing capabilities and adding capabilities outside the firm's previous repertoire. While both options are available to firms, the second one has received less attention and is the one on which this paper focuses.

Design, at the bottom of the wheel of fortune, originally referred to a bespoke process of strategy design directed by the CEO (Andrews 1971) to match a firm's distinctive competences and its external environment. Porter's (1996) work on activity systems elaborates this idea but from a static perspective that emphasizes the linkages across choices at a point in time. Some of his successors have adopted a more dynamic perspective—for example, Siggelkow's (2002) careful description of the evolution of Vanguard's "characteristic features." This paper advances analysis along these lines by examining how a field of strategic possibilities can systematically be searched and expanded, and how optimal strategies for doing so vary over time.

Innovation is often conceived of as searching for components and experimenting with component combinations (e.g., Schumpeter 1939, Stuart and Podolny 1996, Fleming 2001, Arthur 2009). One mathematical structure that strategists have used in recent years to analyze component combinations is NK-modeling of the sort developed by Kauffman. But Kauffman himself notes, in his recent paper with Felin et al. (2014, p. 262), that "the focus on bounded rationality and search [of the NK-approach] is highly problematic for the fields of entrepreneurship and strategy and does not allow us to explain the origins of economic novelty." This paper outlines a new kind of structure for modeling innovation that avoids this problem.

Agility has probably received the least attention from academic strategists, but it has recently attracted substantial interest from practitioners, with concepts such as minimum viable products and lean strategies being emphasized by a broad range of companies (e.g., Doz and Kosonen 2008, Rigby et al. 2016). Arguments for agility are usually tied to a sense of uncertainty that is so overpowering that reacting quickly is the key.[1] However, this literature has not pushed very far on the analytical front. In this paper we make headway by modeling consecutive component choices based on the cumulative information about the unfolding innovation space up to the most recent prior component choice. The range of strategies we derive provides guidance on when an agile strategy is appropriate and when more complex strategies are optimal.

Underpinning this paper's contributions is the idea of leveraging the underlying mathematical structure of a combinatorial innovation process by using the information thrown off by an ongoing innovation search process to create advantaged innovation strategies. This approach has points of contact with several apparently distinct answers to the question of where great strategies come from. And it offers guidance both for policies for component choice in the unfolding innovation process and in the broader sense of how to employ information-enabled search to achieve long-term success.

3. Illustrative Examples

The combinatorial perspective on innovation adopted in this paper can be illustrated with products and the components used to make them in two different sectors. In gastronomy, the products are 56,498 recipes from the databases Allrecipes.com, Epicurious.com, and Menupan.com, and the components are 381 ingredients (Ahn et al. 2011). In technology, the products are 1,158 software products catalogued by StackShare.io, and the components are 993 development tools used to make them.

The *usefulness* of a component is defined in both cases as the number of valid products we can make that contain it. Analysis indicates that the relative usefulness of a component depends on how many other components have already been acquired. For each of the two sectors, Figure 2 displays the usefulness of three typical components, averaged at each stage over all possible choices of the other components that have already been acquired.

One of the key insights from the two product spaces considered concerns *component crossovers*: the relative usefulness of different component building blocks changes over time as the number of components

Figure 2. (Color online) The Relative Usefulness of Different Components Can Cross as the Number of Components Increases in Gastronomy (Left) and Technology (Right)

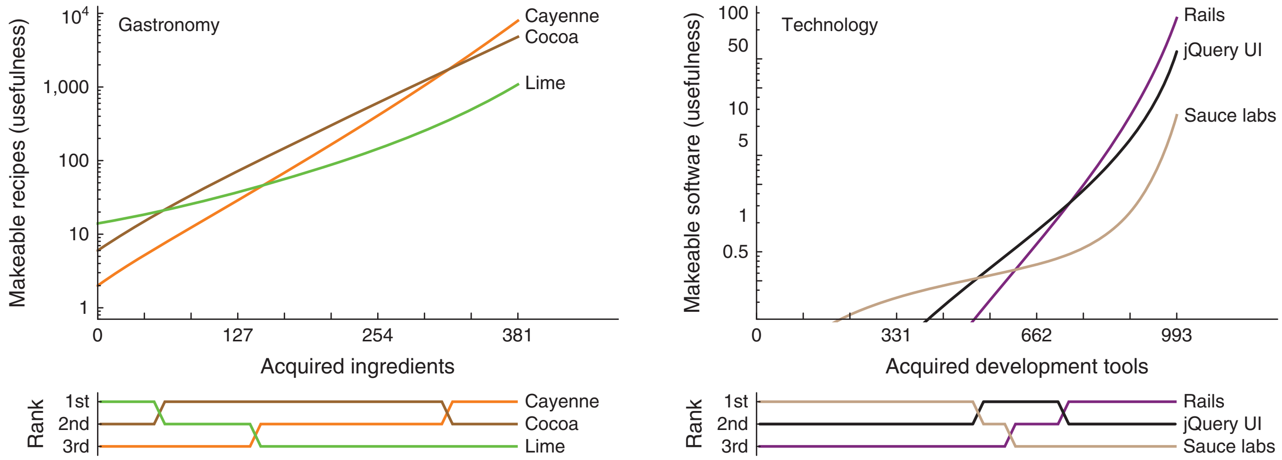

Note. The top panels show the mean usefulness of three typical components, averaged over all possible sets of other components in the baskets for each basket size; the bottom panels show the rank order of components by usefulness. Adapted from Fink et al. (2017).

increases. For example, given a small basket of ingredients, adding cocoa to the basket generally boosts the number of recipes that can be made more than adding cayenne. But the reverse is true with a large basket of ingredients: adding cayenne to the basket is generally more advantageous than adding cocoa. In other words, cayenne tends to be comparatively more useful in making more complex recipes.

For a company-focused example of how to think about innovation from this perspective, consider the case of Apple. Mazzucato (2013) has argued that most of the components that Apple assembled into its breakthrough products were actually the preexisting fruits of state-sponsored research (see Figure 3). Whether one agrees with all the details of Mazzucato's "origin story" or not, the picture that she paints is consistent with the component-based approach adopted here. And more broadly, IBM's Component Business Model, which represents a business as a set of 25–30 building blocks, suggests the utility of a component-based approach for strategy as well as product development.

Within such a setup, adding the most useful component on offer is a good strategy for maximizing the size of the product space immediately. But it is more complicated to make choices now that will most expand the size of the product space in the future. To what extent should an immediate boost to the size of the product space be forgone in order to achieve a bigger, but less certain, boost in the future?

The next section specifies a component-based model of innovation that provides a quantitative means of answering this question. It enables forecasts of the future usefulness of components currently on offer, and it suggests a spectrum of strategies whose optimality depends on, among other things, the resources available to a firm and the maturity of the innovation space.

4. The Model and Analysis

As in the illustrative examples, products (or businesses) are assumed to be made up of distinct components. A component can be a material object or a skill or a routine. Capacity constraints are assumed away: there is more than enough of each type of component for our needs. Any subset of the components that have already been acquired can be combined, but a combination either is—or is not—a valid product, according to some universal recipe book of products. Suppose further that there are a total of N possible components in "God's own cupboard," but that, at any given stage, a firm's basket of components contains only n of these N possible building blocks. At each stage, the firm picks a new component to add to its basket.[2]

So there are two essential elements to the model. First, there is a basket of different kinds of components that grows with time as new components are added to it. Second, there is a prespecified set of valid and invalid combinations of these components, with the valid combinations representing viable products. We assume that the firm knows which combinations of the components already in its basket, as well as combinations from the existing basket with any single new component under consideration for adoption, constitute viable products. But, crucially, we assume no knowledge of the "recipe book" beyond this—that is, we do not know whether or not combinations containing multiple components outside of our basket are viable.

The firm's goal is to maximize the number of products that it can make—its product space—as it adds more components to its basket. The model does not consider the different values associated with specific products, which will depend on the market environment and may change over time. Instead, it simply seeks to maximize the *number* of viable products that it can make,

Figure 3. Origins of Apple Products

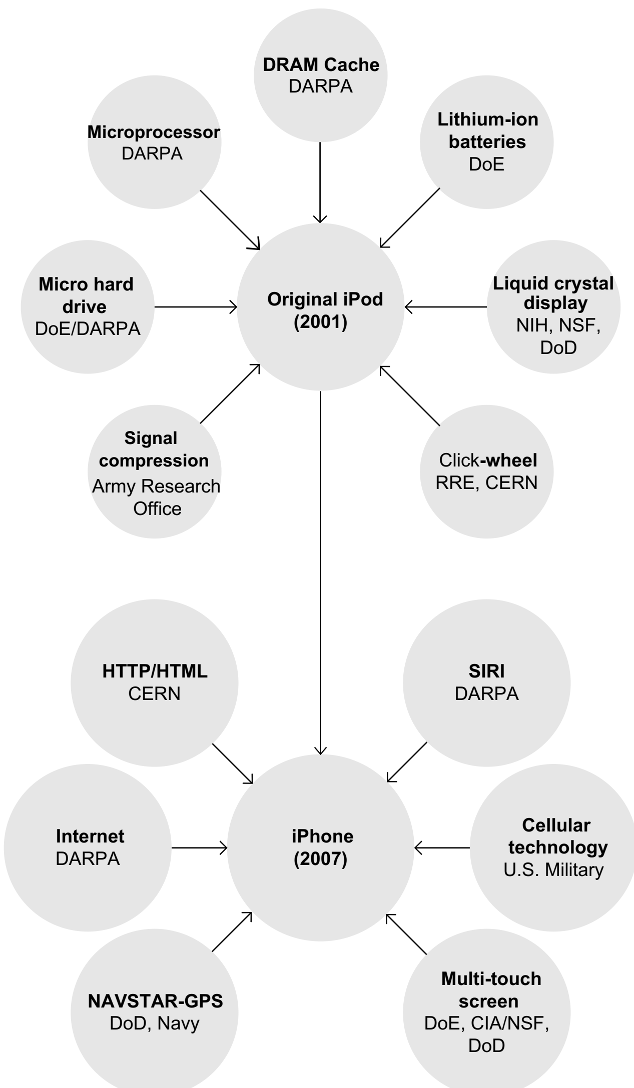

Source. Adapted from Mazzucato (2013) (p. 109, figure 13).

with the knowledge that, on average, some will win and some will flop. In other words, the size of the product space is treated as a reasonable proxy for competitive position; we do not get into the complexities of elaborating a full-fledged product-market subgame.[3] Note that a similar proxy is used in models of evolution, where evolvability is defined as the number of new phenotypes in the adjacent possible, or 1-mutation boundary, of a given phenotype (Wagner 2008).

4.1. Complexity and Usefulness

In line with the illustrative examples, we introduce two simple variables: the complexity of products and the usefulness of components. As we will show, understanding the relationship between these two variables can help us forecast the usefulness of a component in the future given what we know about its usefulness now.

The notions of product complexity and component usefulness are formalized as follows. The complexity c of a product is the number of distinct components it contains. Multiple occurrences of the same component count once, so that the word "apple" is made up of four different component letters, not five. As well as being more mathematically tractable, this is also in keeping with the lack of component capacity constraints that we assumed earlier.

The usefulness u_α of some component α is how many more products we can make with α in our basket than without α in our basket. As we gather more components and n increases, u_α increases or stays the same; it cannot decrease. We therefore write $u_\alpha(\underline{n})$ to indicate this dependence of the usefulness on $\underline{n}$, where $\underline{n}$ is the set of n components in our basket.

This distinction between the basket of components $\underline{n}$ and the number of components in that basket n reminds us that the usefulness of component α depends not just on how many components we have but also on the particular choice of those components. For example, the usefulness of the letter c when our basket contains a and b is greater than the usefulness of c when our basket contains d and e (eight new words versus five).

To sidestep this dependence of usefulness on the particular choice of components, we could, in principle, average over all possible baskets of a given size. But this is not computationally viable for even modest values of n. However, Fink et al. found a mathematical shortcut for calculating this exactly, as described in the methods section of Fink et al. (2017). Armed with this technique, we define the mean usefulness $\bar{u}_\alpha(n)$ in a natural way: the average of the usefulness of α over all choices of the $n-1$ other components in our basket from the $N-1$ possible components. Notice that the mean usefulness $\bar{u}_\alpha$ depends only on the number of components n rather than on the particular choice of those components $\underline{n}$.

To make a product of complexity c, we must possess all c of its distinct components. So making a complex product is harder than making a simple one, because there are more ways that we might be missing a needed component. We therefore group the products we can make containing the component according to their complexity. That is, the mean usefulness $\bar{u}_\alpha(n,c)$ of a component is how many more products of complexity c we can make, on average, with α in our basket than without α in our basket. Summing $\bar{u}_\alpha(n,c)$ over c gives the mean usefulness $\bar{u}_\alpha(n)$, as expected. The advantage of this more refined partitioning is that, by understanding the behavior of $\bar{u}_\alpha(n,c)$, we can understand the more difficult $\bar{u}_\alpha(n)$. Our key result, proved by Fink et al. (2017), can be expressed as a conservation law: for any stage n and some other stage n', $\bar{u}_\alpha(n,c)n/\binom{n}{c} = \bar{u}_\alpha(n',c)n'/\binom{n'}{c}$, where $\binom{n}{c}$ is the binomial coefficient.

When the number of components is big compared with the product complexity ($n, n' \gg c$, which tends to be the case in practice), we can approximate $\binom{n}{c}$ and $\binom{n'}{c}$ by n^c and by n'^c, and therefore

$$\bar{u}_\alpha(n', c) \simeq \bar{u}_\alpha(n, c)(n'/n)^{c-1}. \quad (1)$$

What this tells us is that the usefulness of some component at the present (stage n) is not a good indicator of its usefulness in the future (stage n'). The current usefulness underestimates future usefulness by the factor $(n/n')^{c-1}$, which we see by solving Equation (1) for stage n usefulness. Let us take a closer look at this distorting factor. The more we look into the future, the more it undervalues usefulness. But the distortion does not apply to all products in a uniform way: it gets exponentially worse with complexity, meaning that complex products get undercounted much more than simple products. We call this distortion factor $(n/n')^{c-1}$ the *complexity discount*, to highlight this unusual exponential dependence on complexity.

The correction to the complexity discount is just its inverse—namely, $(n'/n)^{c-1}$: the factor that shows up in Equation (1). Applying this correction to our distorted "vision" of the future based on what we see now, we can make more farsighted choices than we otherwise would. Early on, $\bar{u}_\alpha(n, c)$ will tend to be small for higher complexities, but depending on how far ahead we look, the bigger growth rate can more than compensate for this, as we see in Figure 2. Summing Equation (1) over size c, and noting that u_α is an unbiased estimate of its mean, we find

$$u_\alpha(n') \simeq u_\alpha(n, 1) + u_\alpha(n, 2)x + u_\alpha(n, 3)x^2 \ldots, \quad (2)$$

where $x = n'/n$ is the effective time, with $x = 1$ being the present time. Thus we can think of the future usefulness of component α as a polynomial in time, in which the coefficients are just the current usefulnesses of α for different complexities.

This approach is attractive, in part, because mean usefulness does not depend on the correlations between components—for example, the higher chance of finding onion in recipes that contain garlic. In other words, we do not have to worry about the design structure matrices associated with all of our products. Correlations do influence the size of fluctuations in the usefulness u for a particular basket around the mean usefulness $\bar{u}$, but they do not influence the mean itself.[4]

4.2. Strategies

Our characterization of optimal strategies is based on the complexity discount. Components that tend to show up in complex products do not seem useful early on, because we are likely to be missing other components that those products require. The complexity discount captures the extent to which focusing on the current usefulness of components distorts perceptions of their future value. The key insight is that it is possible to correct for this complexity discount by using currently available information about the products that can be made.

Our research shows that the most important components—materials, skills, and routines—when an organization is less developed tend to be different from when it is more developed. The relative usefulness of components changes over time in a statistically repeatable way. Equation (2) is important because it enables harnessing these crossovers by anticipating them before they happen.

By choosing those components at stage n that maximize the product space at some later stage n', we have a spectrum of strategies, depending on how far ahead we set our sights. A shortsighted strategy (n' close to n) maximizes what a potential new component can do for us now. It considers only the usefulness of a component, because the effective time, x in Equation (2), is close to 1. A farsighted strategy (n' far from n) maximizes what a potential new component could do for us later. It considers both the usefulness of a component and the complexity of products containing it.

A farsighted strategy can outperform an impatient strategy only to the extent that there are component crossovers to forecast. Without crossovers, the two strategies collapse into one. Therefore an important characteristic of any sector is the prevalence of crossovers in it, which can vary substantially. Figure 4 shows the top 40 most useful components as the size of our basket goes from 1 to 381 ingredients (top) and from 1 to 993 technology tools (bottom). Software technology turns out to exhibit more crossovers than gastronomy (see Figure 4).

5. Discussion

The analysis in this paper emphasizes and extends some classic insights about strategy and also develops some new themes.

A key insight from the analysis is that there are different frames of reference for prioritizing a set of components. The most useful components in one frame, or innovation stage, need not be the same as in another. No single frame is inherently more valid than any other; the frame we prioritize depends on our current stage and how far into the future we wish to and are able to look. In broad terms, this frame dependence is reminiscent of special relativity, in which distances, times, and the order of events depend on the inertial frame of reference. In our work, different stage frames are related by the transformation in Equation (2), whereas in special relativity, different inertial frames are related by the Lorentz transformation. In our work, $\bar{u}_\alpha(n, c)n/\binom{n}{c}$ is invariant, whereas in special relativity, the speed of light is invariant. It does

Figure 4. (Color online) The Dependence of the Relative Usefulness of Different Components on the Number of Components Acquired in Gastronomy (Top) and Technology (Bottom)

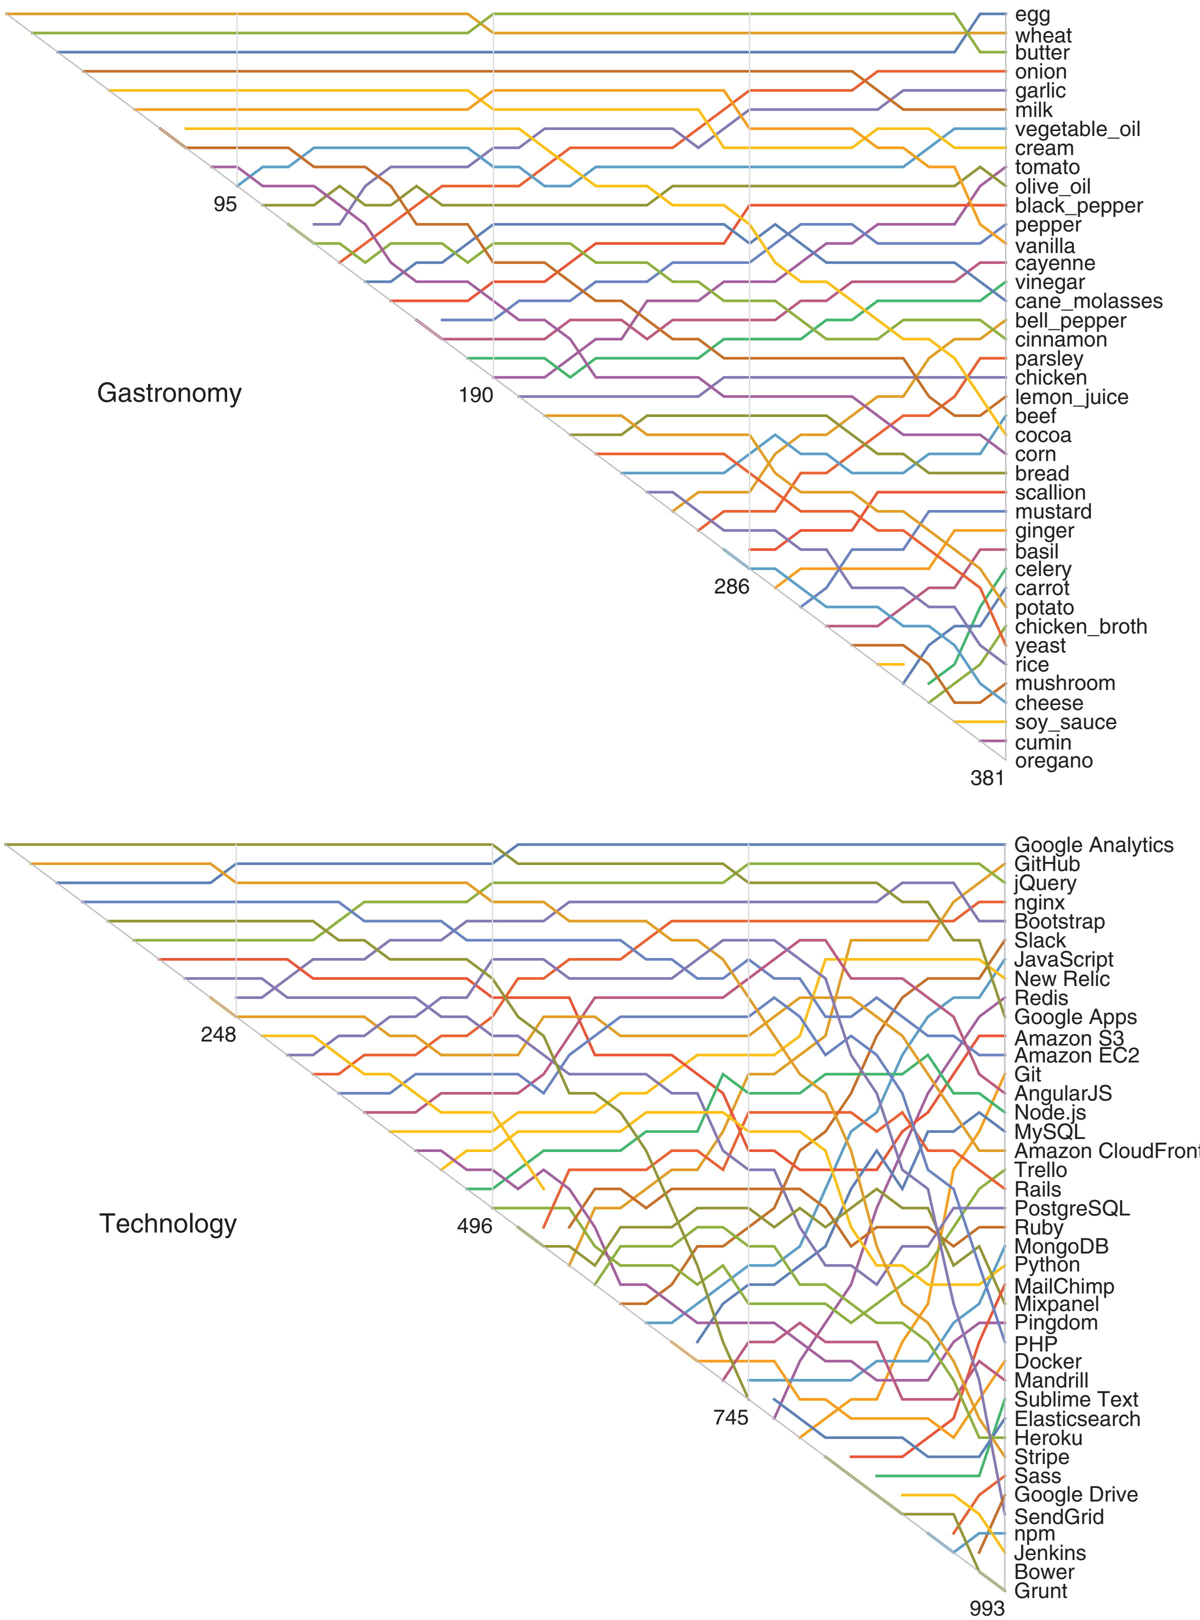

not make sense to ask what the most important components are without specifying a frame: the number of components already acquired. This runs counter to intuition, or at least convenience: we typically regard as absolute the most important ingredients in a kitchen, or tools for survival, or skills in a company.

A second key insight is related to contingency, one of the oldest themes in the strategy literature. Our strategic approach to component acquisition is contingent on factors internal to a firm as well as external to it. Internally, the optimal strategy depends on resource constraints and, more broadly, the objectives of the firm, which are related to its governance.[5] Resource-constrained firms tend to favor an impatient strategy and immediately reap the value of new components, whereas wealthier firms likely favor a farsighted strategy and, after a stagnant period assembling needed components, expect to achieve greater growth as the value of those components kicks in. A similar contrast is to be expected between, on the one hand, businesses overseen by private-equity firms that, given the finite duration of limited partnerships, typically seek to complete their buy-to-sell cycle over a several-year period and, on the other hand, family-owned firms, which are typically supposed to be more interested in long-term growth and resilience than immediate value maximization. And looking externally, the decision to adopt a farsighted strategy depends on aspects of the sector that are beyond the firm's control—namely, the extent to which it exhibits crossovers in component value. A farsighted strategy outperforms an impatient one only to the extent that there are important component crossovers that can be anticipated and exploited.[6]

Third, our analysis stresses the importance of trade-offs, but with a clearly dynamic twist. Trade-offs represent a familiar theme in the strategy literature, but one that has most often been discussed in a static context (e.g., Porter's 1980 work on cost versus differentiation). The analysis in this paper, by contrast, takes a dynamic perspective on trade-offs that has relatively few precedents in at least the analytical branch of the strategy literature. The distinction between impatient and farsighted strategies more closely resembles the distinction in evolutionary biology between r-selection (more offspring) and K-selection (better offspring). The r-selection approach, similar to our impatient strategy, invests little in nurturing individual progeny, focusing instead on fast, immediate growth.

Fourth, our analysis fits with recent academic and practical work on strategy dynamics emphasizing the importance of both irreversibility and uncertainty for dynamic thinking about strategy to really be required (Ghemawat 2016). In regard to the former, Arrow (1964, 1968) noted half a century ago that without irreversibility of any sort, choices could be reversed costlessly and therefore be made myopically, without penalty. Ghemawat (1991) applied these ideas to dynamic strategy. Our model exhibits irreversibility but of a particular sort. The extent to which a new component increases our product space is critically dependent on which components we already have in our basket. This is an example of what Page (2006) has called "phat-dependence," to emphasize that what matter are the components that are available at a point in time rather than the order in which they were acquired. In other words, we assume, as is common in economics and control theory (but not necessarily business strategy), state dependence rather than full-blown path dependence. Despite this weaker form of path dependence, recognizing the role of irreversibility proves to be critically important in our model to the formulation of a successful strategy. Every choice of component impacts not only our immediate product space—in how it combines with the components we already have—but also, more importantly, our product space into the future—by combining as well with the components we have yet to acquire. The second condition for dynamic thinking to really be required, uncertainty, is inherent in our model through the assumed lack of knowledge of which combinations of components outside of those in our basket form products. It is this that prevents us from computing the optimal strategy at any point in time with certainty. This theme is particularly worth highlighting given the strong bias toward underestimating uncertainty and overestimating the ability to control outcomes. A Boston Consulting Group survey of 120 companies in 10 industry sectors suggests a strong intellectual and behavioral bias toward classical strategic planning, even in environments where a more adaptive approach would seem optimal (Reeves et al. 2012).

Our approach is particularly suited to adaptive navigation through large, dynamic search spaces, where simple deductive frameworks are likely to be inadequate. As such, this approach builds a bridge between cross-sectional work on strategy that is preoccupied with high dimensionality (e.g., Porter's 1996 detailed activity systems) with dynamic perspectives that typically posit few dimensions because of the difficulty in navigating high-dimensional landscapes. By eschewing deductive models and grounding our choices on information accumulated up to the present moment, we are able to provide a quantitative model for innovation while acknowledging the complex and shifting environment in which real companies operate.

We now return to the six origin stories that we briefly outlined in Section 2, adding to each some specific contributions from the analysis in this paper.

Without discounting the role of *luck*, our understanding of component crossovers offers new insight into serendipity: the fortunate but unforeseen events that many firms identify as key to some of their best products. Components that depend on the presence of many others can be of little benefit early on. But as the innovation process unfolds and these dependent components pay off, the results will seem serendipitous, because a number of previously low-value components become more valuable. In this context, what appears to be serendipity is actually the delayed fruition of components reliant on the presence of others. Also note that while the word "serendipity" does not have a well-established antonym (though "zemblanity," which means foreseeable misfortune, comes close), for every beneficial shift in a crossover, there is a detrimental one. Each opportunity for serendipity goes hand-in-hand with a chance for "antiserendipity": the acquisition of components useful now but less useful later. Avoiding these overvalued components can be as important as acquiring undervalued ones in order to secure a large future product space.

In terms of superior *foresight*, the analysis suggests ways of correcting for the complexity discount by using information about the products that we can already make, to make farsighted component choices. Many innovators already have the tools to do so. Companies routinely reengineer competitors' products, analyze the patent landscape, and conduct interviews with technology experts to guide their operational decisions. At the same time, there is an ongoing explosion of digital information about products, competitors, and users. Companies can use these tools and information to guide their strategy by methodically measuring the evolution of product complexity in their space. This requires, in the first instance, the development of a taxonomy of not only physical components but also intangible ones such as skills and routines.

Capabilities, like material objects and processes, can be seen as a particular manifestation of components. Our approach reinforces the importance recently stressed by Pisano (2017) but inherent to Penrose's (1959) original conception of firms' opportunity sets expanding over time, of thinking about broadening capabilities as well as the more customary focus on deepening them. While both are important, the digitalization of business arguably expands the range of component choices and increases the interconnectability of those components, leading to an explosion of combinatorial possibilities. A strategic understanding of which components to adopt, and when to adopt them, is more crucial than ever. And while the capabilities literature, including Pisano (2017), has tended to focus on specialization versus fungibility, our analytical approach suggests that some additional characteristics of capabilities, such as the average complexity of products containing them, are key to effective capability development.

In regard to design, we normally take the complete space of products as a given, to be successfully navigated through an appropriately ordered choice of components. But in some instances, the entire product space forms the effective product, and we have the design opportunity to reverse engineer it so as to promote optimal navigation. Consider software, where different commands (components) can be combined in different ways to perform tasks (products). Learning new commands requires effort, and if the amount of return on effort (the number of new tasks that can be achieved after learning a particular command) is not high enough, the user is liable to give up. Thus it is important to design software so that the "product space" of the user grows at a rate sufficient to motivate persistence with learning new "components." More generally, by connecting strategy problems with their underlying information problems, our approach represents one step toward many product-based businesses reconceiving themselves as information businesses.

Our model of *innovation* is fundamentally component based, but the word "component" belies the generality of the concept. A component can be a physical building block, such as the multitouch screen in Figure 3, but it can also be a skill, such as advertising on social media, or a routine, such as a client survey. Our key insight is the subjectivity of any one frame of reference for ordering component building blocks, and we elaborate on this and its consequences elsewhere. This applies as much to business processes and strategic frameworks as it does to new technologies.

Finally, as far as *agility* is concerned, our approach, with its emphasis on the dependence of an organization's priorities on its maturity, provides a contingent perspective on whether agile/lean approaches make sense. Our analysis supports such approaches to building companies and launching products when resources are in short supply and time horizons are short. Without the resources to sustain a farsighted approach, many start-ups do need to bring a simple product to market very quickly. On the other hand, firms that can weather an initial dearth of new products are likely to see their sacrifice pay off when their farsighted approach kicks in.

6. Conclusion

The editors of the special issue invited us to adopt a prescriptive approach to identifying where great strategies come from. What would we tell our students in response to this question? There is a broad range of answers, some of which we outlined in the introductory wheel of fortune. While they have some overlap, the six origin stories make different assumptions and offer markedly diverse insights. This paper takes a step toward developing a unifying perspective, starting from the context of a component model of innovation. Our main insight is that the most important objects, skills, and routines are not static but depend on how far along the innovation process a firm has progressed. By quantifying the relative usefulness of components at different stages, our analysis helps interpret aspects of luck, foresight, and agility on the one hand and design, capabilities, and innovation on the other, in the context of the same underlying mathematical structure.

Our model is amenable to additional elaboration. Consider three directions that we have not examined in detail but that seem worthy of additional exploration. First, in his classic account of the evolution of technology, Arthur (2009) adopts a component model of technology, but one that allows recursion: every product can be used as a component in making new products. Incorporating recursion into our framework affords a basis for reconceptualizing the popular idea of disruptive innovation, in which the complexity of swathes of product space can be reduced by a single substitution. Second, in practice, adopting each new component has a cost to be offset against the payoff from the expanding space of products. Making this exchange explicit provides a model for quantifying poverty traps: the self-reinforcing mechanism in which governments invest in inferior capabilities with a limited, short-term payoff. Without the resources to weather the delayed payoff of a farsighted strategy, this can give rise to a vicious circle of decline. Third, and most ambitious, is the injection of competition into the picture. With multiple firms each competing for not only large but also distinct product spaces, learning about others' product spaces can be as important as forecasting one's own. This opens up consideration of concepts such as imitation and cooperation.

Acknowledgments

The authors have benefited from comments by Bruno Cassiman, Deepak Hegde, and the editors of the special issue on an earlier draft of this paper.

Endnotes

[1] Of course, with even more uncertainty, we would come full circle, back to luck.

[2] While we at no point posit knowledge of N, the number of possible components, we take it to be much larger than n, the number of components we have acquired. This means that we are at no point picking among the dregs of dead-end components. This is key to the validity of our conclusions.

[3] Theoretical predictions regarding product market subgames in the presence of product differentiation tend to be very sensitive to the modeling assumptions employed (e.g., Anderson et al. 1992). But with Bertrand competition in prices and less than complete coverage of product space (i.e., fewer products on offer, overall, than in the universal recipe book), the firm capable of making more products should expect to earn higher operating profits—although even this conclusion abstracts away from prior strategizing, in an interactive sense, about product (component) selection and introduction.

[4] The predictive capability of our method of forecast rests on the assumption that the fluctuations around the mean are small compared with some relevant change in the mean itself. We have measured this for specific gastronomy components, and it tends to be valid.

[5] While this point has gotten obscured by the recent emphasis in strategy on shareholder value maximization, Andrews (1971), one of the founding fathers of the strategy field, noted to one of us (Ghemawat) 20 years ago how surprised he was that the concept of strategy had gotten divorced from consideration of purpose.

[6] Although we do not consider this possibility here, the ability of a farsighted strategy to outperform may also depend on the maturity of the innovation space in sectors in which the number of possible components is sufficiently limited to be exhausted.

References

Ahn YY, Ahnert SE, Bagrow JP, Barabási AL (2011) Flavor network and the principles of food pairing. *Sci. Rep.* 1:Article 196.

Anderson SP, de Palma A, Thisse JF (1992) *Discrete Choice Theory of Product Differentiation* (MIT Press, Cambridge, MA).

Andrews KR (1971) *The Concept of Corporate Strategy* (Dow Jones-Irwin, New York).

Arrow KJ (1964) Optimal capital policy, the cost of capital, and myopic decision rules. *Ann. Inst. Statist. Math.* 16(1):21–30.

Arrow KJ (1968) Optimal capital policy with irreversible investment. Wolfe JN, ed. *Value, Capital and Growth* (Edinburgh University Press, Edinburgh, UK), 1–19.

Arthur WB (2009) *The Nature of Technology: What It Is and How It Evolves* (Simon and Schuster, New York).

Barney JB (1986) Organizational culture: Can it be a source of sustained competitive advantage? *Acad. Management Rev.* 11(3):656–665.

Caves RE, Gale BT, Porter ME (1977) Interfirm profitability differences: Comment. *Quart. J. Econom.* 91(4):667–675.

Doz Y, Kosonen M (2008) *Fast Strategy: How Strategic Agility Will Help You Ahead of the Game* (Pearson/Longman, Harlow, UK).

Felin T, Kauffman S, Koppl R, Longo G (2014) Economic opportunity and evolution: Beyond landscapes and bounded rationality. *Strategic Entrepreneurship J.* 8(4):269–282.

Fink TMA, Reeves M, Palma R, Farr RS (2017) Serendipity and strategy in rapid innovation. *Nature Comm.* In press.

Fleming L (2001) Recombinant uncertainty in technological search. *Management Sci.* 47(1):117–132.

Ghemawat P (1991) *Commitment: The Dynamic of Strategy* (Free Press, New York).

Ghemawat P (2016) Evolving ideas about business strategy. *Bus. Hist. Rev.* 90(4):727–749.

Mazzucato M (2013) *The Entrepreneurial State: Debunking Public vs. Private Sector Myths* (Anthem Press, London).

Page SE (2006) Path dependence. *Quart. J. Political Sci.* 1(1):87–115.

Penrose E (1959) *The Theory of the Growth of the Firm* (Oxford University Press, Oxford, UK).

Pisano G (2017) Towards a prescriptive theory of dynamic capabilities: Connecting strategic choice, learning, and competition. *Indust. Corporate Change* 26(5):747–762.

Porter ME (1980) *Competitive Strategy: Techniques for Analyzing Industries and Competitors* (Free Press, New York).

Porter ME (1996) What is strategy? *Harvard Bus. Rev.* 74(6):61–78.

Reeves M, Love C, Tillmanns P (2012) Your strategy needs a strategy. *Harvard Bus. Rev.* 90(9):76–83.

Rigby DK, Sutherland J, Takeuchi H (2016) Embracing agile. *Harvard Bus. Rev.* 94(5):40–50.

Schumpeter J (1939) *Business Cycles* (McGraw-Hill Book Company, New York).

Siggelkow N (2002) Evolution toward fit. *Admin. Sci. Quart.* 47(1):125–159.

Stuart TE, Podolny JM (1996) Local search and the evolution of technological capabilities. *Strategic Management J.* 17(S1):21–38.

Wagner A (2008) Robustness and evolvability: A paradox resolved. *Proc. Roy. Soc. B* 275(1630):91–100.

Thomas Fink is director of the London Institute for Mathematical Sciences and Charge de Recherche in the CNRS. He uses statistical physics and discrete mathematics to form basic science insights into complex physical, biological and socioeconomic systems. His research interests include evolvability and technological innovation, resilience and antifragility, graph theory, enumerative combinatorics, and fractal mechanics. He studied physics at the California Institute of Technology and Cambridge University, winning the Fisher Prize for physics and Green Prize for research. He did a postdoc at École Normale Supérieure and was a junior fellow at Caius College, Cambridge, where he worked at the Cavendish Laboratory. Outside of physics, he is interested in models of simplicity, collective creativity, graphic and information design, and mathematics in culture.

Pankaj Ghemawat is Global Professor of Management and Strategy and director of the Center for the Globalization of Education and Management at the Stern School of Business at New York University, and the Anselmo Rubiralta Professor of Global Strategy at IESE Business School. His research focuses on sustainable competitive advantage, strategic investment, globalization, and emerging markets. A fellow of the Strategic Management Society and the Academy of International Business, he is also the founding chairman of the Foundation for Practice and Research in Strategic Management (PRISM). He graduated from Harvard College with a degree in applied mathematics and economics and received his Ph.D. in business economics from Harvard University.

Martin Reeves is director of the BCG Henderson Institute, the Boston Consulting Group's thinktank for new ideas in strategy and management, and senior partner in BCG's New York office. His research interests include intervening in complex adaptive systems, biological approaches to management, restoring and maintaining corporate vitality, gaming and strategy, intelligent organization, and the science of transformation. He graduated from Cambridge University in natural sciences, studied biophysics at Tokyo University, and has an MBA from Cranfield University.