

The Dynamics of Capability Development and Erosion

Hazhir Rahmandad

Assistant Professor, Industrial and Systems Engineering
Virginia Tech,
Northern Virginia Center- Rm 430
7054 Haycock Road, Falls Church, VA 22043
703-538-8434
hazhir@vt.edu

Nelson Repping

Associate Professor, Sloan School of
Management, M.I.T
E53-335, 30 Wadsworth Ave.,
Cambridge, MA 02142
617-258-6889; nelson@mit.edu

The Dynamics of Capability Development and Erosion

Abstract

While the notion of a capability is widely invoked to explain differences in organizational performance, we know little about where capabilities come from or why some organizations retain higher levels of a capability than others. To develop an improved understanding of the micro-foundations of capability, we study two software development organizations that showed significantly different capability trajectories despite similar organizational and technological settings. Our analysis of these data suggest the existence of an organizational dynamic that we call the adaptation trap: due to a failure to account fully for delays in receiving information feedback, well-intentioned efforts by managers to search locally for the optimal resource/workload balance lead them to systematically overload their organizations and, thereby, cause capability to erode. We provide a formal model of our emerging theory, the analysis of which provides both insight into the sources of differential organizational performance and testable hypotheses.

Acknowledgments:

Special thanks to Martine Devos, David Weiss, Randy Hackbarth, John Palframan, and Audris Mockus and several members of Sigma organization for their support of, and contribution to, this research. We are also grateful for the valuable feedback of Michael Cusumano, Rebecca Henderson, Shayne Gary, Jan Rivkin, John Sterman, Jeroen Struben, Lourdes Sosa, Gokhan Dogan, Ezra Zuckerman and several seminar participants at MIT, Stanford, Wharton, and the Academy of Management.

Introduction

The field of strategy has long concerned itself with understanding why some firms do better than others. While arguments focused on position (Porter, 1998) continue to offer both significant explanatory power and insight to practitioners, firms that manage to produce above-average results without the benefit of an obvious positional advantage have increasingly become the focus of scholarly attention. Inspired by the seminal work of Nelson and Winter (1982), Wernerfelt (1984), and Barney (1991), a growing number of scholars now look inside the firm for the sources of competitive advantage and heterogeneous performance.

As befits the diversity of the perspectives that characterize organization theory, the search for competitive advantage within the firm has been guided by several theoretical and methodological frames. Scholars favoring large-sample approaches have produced several analyses suggesting that heterogeneity in firm performance cannot always be fully explained by positional differences (Rumelt, 1991; Mauri, and Michaels, 1998; Spanos, and Lioukas, 2001). This line of work suggests both that what happens inside the firm matters and that differences in internal operations are not quickly eliminated by imitation (Henderson, and Cockburn, 1994; Roberts, and Dowling, 2002). Building on these studies, theorists have pursued a number of paths towards a convincing explanation of the observed differences in performance. The literature now includes a growing collection of formal models that demonstrate the difficulties associated with imitating successful firms (e.g. Rivkin, 2001), a set of studies detailing both the development of internal resources and assets (Siggelkow, 2002) and, another set focusing on the failure of firms to adapt those resources to changing environmental requirements (e.g. Henderson, and Clark, 1990).

A central notion arising from this inquiry is the *capability*, a set of routines that enable a firm to produce a particular output given a set of inputs and targets (e.g., develop new products) (Winter, 2003). A capability differs from the normal conception of an asset in that it cannot be purchased directly but instead is the accumulated product of a particular set of activities (e.g. Dierickx, and Cool, 1989; Winter, 2000). Positing capability as the product of accumulation creates several platforms from which to build a theory of persistent heterogeneity in firm performance. Most fundamentally, because capabilities in this conception cannot be acquired in external markets, one of the most powerful forces for competitive equilibration is eliminated. Instead, firms wishing to close a competitive gap must resort to either imitation—growing their own capabilities to match those of a competitor—or substitution (Dierickx, and Cool, 1989; Peteraf, 1993).

Both paths to strategic parity are, however, fraught with pitfalls. Imitating the processes through which a particular firm developed a given capability is impeded by the complexity of modern organizations (Rivkin, 2000; Rivkin, 2001). For example, despite having studied for more than two decades, the question of which elements of the Toyota production system are responsible for Toyota's manufacturing capability remains much debated and few substitutes have emerged for Toyota's unique abilities. By positing the existence of a transformation process that cannot be purchased directly, the notion of capability provides a platform from which to resolve one of the key questions facing organizational scholars: why do seemingly similar firms produce significantly different levels of performance?

Despite its promise, the explanatory power of theory built on the existence of capabilities remains modest, and the lack of progress can be traced to the fact that we know little about the

specific mechanisms through which capabilities are created and sustained (Helfat, 2000; Helfat, and Peteraf, 2003). Building on Nelson and Winter (1982), several theorists have advanced the notion that capabilities result from the repeated execution of specific actions and experiential learning (Zollo, and Winter, 2002). Given the breadth of organizational action that stems from routines, however, this view offers little analytical purchase; current theory offers few criteria on which to judge whether a given set of activities and learning mechanisms yield a particular capability. Consequently, the current literature lacks a general explanation for why some firms accumulate capability faster than others. Until such a theory exists, the full potential of the capability concept in explaining the sources of heterogeneous firm performance is unlikely to be realized (Williamson, 1999; Priem, and Butler, 2001).

Capability Dynamics and Dynamic Capabilities

The limitations of the current theory concerning the origin of capabilities can be traced to three sources, conceptual ambiguity, operational distance, and research design.

Conceptual Ambiguity. The discourse on capabilities lacks conceptual precision on a key dimension. Early critics of the resource-based view challenged its explanatory power by arguing that resources central to firm performance could, in principle, be purchased in external markets, thus eliminating them as a source of differential competitive advantage. Following the conceptualization of Dierix and Cool (1989), scholars responded by positing strategically important resources, soon to be called capabilities, as entities that, by definition, could not be purchased, but could only be accumulated through doing particular sets of activities within the firm. Such accumulation processes could, this view held, result in the ability to execute certain transformation processes in ways that would be difficult for other firms to duplicate.

Subsequent scholars labeled those capabilities focused on transforming inputs into outputs as zero-level or operational capabilities, reserving the label “dynamic” for the ability to modify operational capabilities (Teece, et al., 1997; Eisenhardt, and Martin, 2000; Winter, 2003). The discussion then focused on heterogeneity in dynamic capabilities as the main driver of heterogeneity in operational capabilities and thus firm performance. The result is that with a few notable exceptions (e.g. Helfat, and Peteraf, 2003), operational capabilities have been treated as entities that can be reconfigured quickly through dynamic capabilities (Schreyogg, and Kliesch-Eberl, 2007). Such a framing, however, directly contradicts the conceptualization of capabilities as the product of accumulation; accumulation is, by definition, a dynamic process (Sternan, 2000). Moreover, major shifts in capabilities appear to require significant periods of time (Argote, and Epple, 1990; Sternan, et al., 1997; Repping, and Sternan, 2002). Thus, a conceptualization of capability capable of explaining firm heterogeneity must account for the empirical observation that capability can both accumulate and erode. While the available evidence suggests that capabilities can, with the passage of time, be both gained and lost, existing theory offers little to explain how or why.

Operational Distance. The second limitation of existing theory lies in its distance from the underlying phenomenon. As argued by several scholars, the study of, and theory about, organizations has become increasingly removed from the day-to-day activities that constitute organizational life (e.g. Barley, and Kunda, 2001; Kaplan, and Henderson, 2005). Such a critique is particularly relevant to the study of capability. While some theoretical frameworks capture the learning processes underlying capability evolution (Zollo, and Winter, 2002; Gavetti, 2005), empirical work establishing the micro-foundations of such frameworks is virtually non-existent.

Capabilities are not instantiated in the boardroom, but instead reside on production lines, in R&D labs and at the point of customer contact. Few scholars have, however, systematically observed the locations at which capabilities are said to reside.

Such operational distance limits the continued growth of our understanding of capabilities and their connections to firm-level heterogeneity. Without a clear mapping between theoretical concepts and concrete observations, scholars will be hard pressed to confront theories premised on capabilities with data. Worse, without a clearly specified set of mechanisms grounded in observable features of organizational life, empirical research focused on capabilities risks being tautological; capabilities are those things which good firms have rather than being the true source of above-average performance (Williamson, 1999; Priem, and Butler, 2001).

Research Design. Many of the limitations of current theory can be traced to the design of existing research. Of the few studies that attempt to observe capabilities in action, most focus on how firms react to radical technological change (e.g. Henderson, and Clark, 1990; Rosenbloom, 2000; Tripsas, and Gavetti, 2000; Keil, 2004) ((Raff, 2000) is a notable exception). Such a research strategy risks conflating two conceptually distinct determinants of firm performance, the ability to react to change and the ability to execute a distinct set of policies to produce capabilities. While strategy scholars have made substantial progress in understanding how existing capabilities constrain a firm's ability to react to changes in its competitive environment (Leonard-Barton, 1992), we know little about how those capabilities got there in the first place (Gavetti, 2005; Lavie, 2006). Moreover, while scholars have recognized that the current level of capability is determined by both accumulation and erosion (Helfat, and Peteraf, 2003), we know little about the micro-processes underlying the erosion of capabilities. Thus, the research designs

that have dominated the study of capabilities have created a potential blindspot in our understanding of the sources of competitive advantage. Until we understand the micro-processes through which capabilities are built, maintained, and lost, scholars may find it difficult to explain differences in performance that do not stem primarily from an advantaged market position.

Studying the Micro-Foundations of Capabilities

Studying the origin of capabilities poses several challenges. First, capabilities are embodied in the adherence to, and departures from, standard procedures, and are often not things that can be counted, thus making it difficult to identify those activities that contribute to competitive advantage. As much research has demonstrated, such situations are prone to spurious causal inferences and consequent superstitious learning—a situation that afflicts researchers and practitioners equally. Three research design elements are necessary to offset this challenge. First, following the discussion above, capabilities, even operational ones, must be studied longitudinally. If, capabilities are accumulated through a particular set of activities, then their origins can only be understood by explicitly focusing on the accumulation and erosion processes. Second, a comparative design is essential to isolate those practices that actually differ across organizations that display different levels of performance. Cross-sectional comparison, while helpful, will not eliminate the potential for causal ambiguity; higher performing firms may differentially adopt irrelevant practices. Thus, the third necessary element is a clearly specified and testable causal theory that connects more primitive features of the organization in question to the capabilities that are hypothesized to generate differences in performance. Such a theory would allow analysts to treat differences in industry structure and technology as a natural experiment, the analysis of which might yield testable predictions concerning the capabilities

that are both valuable and hard to build and sustain within a specific industry structure. For example, a theory indicating that senior executive team flexibility was a relatively more important and harder to sustain capability in rapidly changing markets would then predict that the return to flexibility is relatively larger in such markets than in those with a slower rate of change. A study designed to *develop* new theory cannot, on its own, eliminate the chance of spurious inference (since it is by definition sampling on the dependent variable). But, by specifying a logically consistent theory, such a design enables the derivation of testable hypotheses and thus provides a check on its utility.

The necessity of a comparative design creates an additional challenge because the level of capability and its rate of accumulation are influenced by innumerable variables beyond the practices that instantiate them. The relative importance of a given capability varies across industries—reliability is more important in airframe construction than in word processing software—and thus firms are likely to focus on them differently. Focusing within a specific industry segment and technology reduces this problem, yet firms within a given industry may have different positions and, therefore, face different constraints in their ability to invest in different capabilities, thus creating an additional source of variability.

The ideal research design would thus allow one to compare the accumulation paths of an important capability across two different organizations that displayed significant differences in performance while holding constant as many outside variable as possible. From such a comparison one could then induce a theory detailing the practices that contribute to the successful accumulation and maintenance of one or more important organizational capabilities. While no actual design can exactly match this ideal, the company discussed in this paper

presented us with the opportunity to study two organizations engaged in the same basic task, software development, facing virtually identical resource constraints, but who nonetheless displayed significantly different trajectories in their core capability. Comparing these two organizations and their approach to software development yielded a theory of capability development, maintenance and erosion with sufficient specificity that it could be formalized in a mathematical model. The resulting characterization of capability dynamics yields both new insights and testable hypotheses.

DATA AND METHODS

The Research Site

Sigma is a large telecommunication firm that develops a variety of software and hardware products. Sigma employs approximately 16,000 who are distributed in multiple locations around the world. In the last five years the company has shifted its strategy towards the software side of the business. Given this emphasis, efficiently developing high quality software is core to Sigma's ongoing success—it is an important operational capability. Despite its centrality to company's mission, however, there is substantial variation in software practice across Sigma's divisions. Some areas consistently produce high quality code while others are plagued with low productivity and high defect rates. Not surprisingly, this variance is of great concern to Sigma's leadership and they were thus eager to participate in a study of their development efforts.

To isolate the practices that separated high performing areas from low performing ones, we chose a polar types research design (Eisenhardt, 1989). With Sigma's help, we identified Alpha and Beta, two products that showed significant differences in performance outcomes despite comparable technical complexity, resource loading, and incentive structures. Both are large

software systems with multiple releases, follow similar official development processes and have similar incentive structures. Despite these similarities, and the initially successful market performance of both products, the processes used by the two development teams have diverged. Beta's practice continues to follow Sigma's suggested practice closely. Alpha's practice, in contrast, has changed significantly, moving progressively towards a "firefighting" mode (Repenning 2001) in which its development resources are increasingly focused on fixing problems rather than the activities that would have prevented those problems.

The two projects also differed in terms of quality, timeliness, and productivity. The Beta team had a history of releasing high quality software on time with few bugs and customer complaints. Alpha's initial releases were also highly successful, but its subsequent performance deteriorated significantly. The comparison of these two projects thus provided us with an opportunity to study how organizational capabilities evolve over time and converge to markedly different states despite similar exogenous conditions and the absence of technological shocks. Comparing and contrasting what happened in the two cases provided a rich context to build a model linking specific practices to the ongoing evolution of an important capability.

Data Collection

The data collection began with the first author spending three months on site at Sigma. During that time he conducted 48 initial semi-structured interviews, each of which was done either in person or on the phone and lasted between 30 and 90 minutes. The initial interviews included members of all functional areas involved in development, services, and sales, architects and systems engineers, developers, testers, customer service staff, sales support personnel, and marketing personnel, as well as managers up to two layers in several of these areas. During the

initial three month period, the first author also attended group meetings focused on the products under study and conducted an extensive review of the available archival data. Later, as the key themes emerged from the analysis, additional interviews were conducted to evaluate and expand emerging hypotheses. These efforts included 22 additional interviews and a group session that elicited ideas from 26 experienced members of the organization. The first author also visited two other sites where work on Alpha and Beta was being conducted.

The actual processes used to develop software (as opposed to that captured in the Sigma manual) is composed of the routines and operating procedures used on a day-to-day basis. Much of our effort focused on uncovering this process and clearly separating what was actually done from what Sigma prescribed. The archival data, mostly from Alpha's case, included project planning and review files for different releases of the product, internal project assessments conducted at the end of major projects, and quantitative data on code development and bug fixing, human resources, and problems reported by customers.

Analysis

We analyzed the data in two phases. The first phase followed standard practice for qualitative research (e.g. Glaser, and Strauss, 1967) and focused on identifying emerging categories and connecting causal hypotheses. In the second phase we integrated these hypotheses into a single cohesive theory via the development of a formal model. Although we describe these two phases in sequence, they overlapped significantly. Attempting to capture our emerging theory in a formal model often necessitated a return to the source data to answer questions raised by the modeling that had not been part of the initial hypothesis generation.

Qualitative Hypothesis Generation. Interviews were recorded and a summary file of each interview was created soon after the interview session. In each interview we searched for factors connected to the practice of developing software and how those factors might relate to major performance metrics (quality and schedule in this case) and looked for corroborations or disagreements with old factors. Based on these factors and the observations on the site, we generated several hypotheses about the processes that contributed to the observed differences in the capability levels and performance outcomes across multiple releases of Alpha and Beta. The complete set of these dynamic hypotheses is documented in the online technical appendix. We used this information to build a conceptual model of how the mechanics of the development process work. We then narrowed the list of hypotheses into a core set based on the themes that were most prominent in the interviews and in the history of the products.

Formal Theory Development. The set of dynamic hypotheses that emerged from our data constituted a qualitative theory explaining the differing performances of Alpha and Beta. While such theories are widely offered in the organizational literature, a variety of studies show that our ability to draw correct inferences from even the simplest dynamic mechanism is limited (See Sterman, 1994 for a summary). To offset these limitations, we used our qualitative analysis as the basis for constructing a formal model to explain the dynamics we observed (Forrester, 1961; Sterman, 2000; Davis, et al., 2007). We used archival data and the judgment of organizational participants to parameterize the model. The modeling process included first building a detailed model validated by the historical performance data and then extracting the main dynamics in the form of a simpler, more stylized model that captures the essential elements of our emerging theory. This simple, generic model is discussed in detail in this paper (a full report of the detailed model can be found in (Rahmandad, 2005)).

The formal model served two purposes in our effort to identify and understand the sources of capability. First, by flagging missing and inconsistent components in our initial set of hypotheses, the formal model enforced the internal consistency of our theory. Iteratively moving between model and data highlighted the pressures and misperceptions that led to the persistence of different decision-rules in the two projects. These decision-rules combine with the physical constraints inherent in developing software to explain the differential accumulation of development capability between Alpha and Beta and their different performance levels. Second, via sensitivity analysis and other simulation experiments, the model enabled the derivation of testable predictions that extend beyond general statements linking capabilities to outcomes and identify the contextual variables that determine the relative value of the capabilities we study. While not completely eliminating the possibility of spurious inference, such specificity positions future researchers to provide more compelling empirical tests of the existence and contribution of the capabilities hypothesized to create competitive advantage.

VARIANCE IN SOFTWARE DEVELOPMENT AT SIGMA

The Espoused Development Process at Sigma

We first outline Sigma's prescribed software development process, which all projects were supposed to follow, and then elaborate on the differing performance of the two projects we studied. The practices that constitute an effective software development process are now well-established (Jones, 2000), and Sigma's proscribed process followed accepted best practice for waterfall development closely. Despite this acceptance, many projects at Sigma, including Alpha, found it difficult to follow these practices on a regular basis.

Software is typically developed in three main phases, which can be followed serially (i.e. waterfall development, typically for larger and well-defined projects) or iteratively (for smaller projects or those with unclear requirements). During the period under study, Sigma's development process was largely serial with the occasional iterative element. In the concept design phase, the main features to be added to the next release are chosen, the product architecture is designed, and general requirements for developing different pieces of code are established. For example, a new feature for the next release of a call-center system could be "linking to national database of households and retrieving customer information as the call is routed to a customer service agent." Software architects determine the method by which this new feature should be incorporated into the current code, which new modules should be written, and how the new code should interact with the product. Subsequently, a more detailed outline of the requirements is developed that describes the inputs, outputs, and performance requirements for the new modules, as well as modifications of the old code.

The next step, usually the largest share of the work, is developing the code. Software engineers (developers) develop the new code based on the requirements they have received from the concept design phase. The quality of the development depends on several factors, including the skills of the developers, complexity of the software, adherence to good practice (e.g., writing detailed requirements and doing code reviews), quality of the code being built upon, and quality of the software architecture (MacCormack, et al., 2003). Preliminary tests are often run in this stage to insure the quality of the new code in isolation before it is integrated with other pieces and sent to the quality assurance stage.

Quality assurance entails testing the code to find problems that stop the software from functioning as desired. Often these tests reveal problems in the code that are referred back to developers for rework. This rework cycle continues until the software passes most of the tests and the product can be released. Note that not all possible errors are discovered during testing. Often tests cover only a fraction of possible combinations of activities for which the software may be used in the field. Therefore, there are always some bugs in the released software.

When the software is released, customers (typically large organizations in Sigma's case) buy and install it and subsequently require service from the company. The service department follows up on bugs reported by customers and refers these bugs to the development organization. Bugs, if they significantly detract from customer satisfaction, are fixed at the customer site through an activity known as *Current Engineering* (CE). CE is often done by the same team that initially developed the code. If the bug is not critical, then it can be fixed in a subsequent release of the software, an activity known as *Bug Fixing*. Bug-fixing usually competes with addition of new features to the next release.

The development, launch, and service activities discussed above focus on a single release (version) of the software, but there are typically multiple releases under development, resulting in two key interconnections across the different releases in progress. First, subsequent generations of software build on their predecessors and hence carry over problems in code and architecture if those problems are not remedied through bug-fixes or refactoring of the architecture (*Refactoring* is improving the design and architecture of the existing code.) Second, because they all draw off a common resource pool of developers, resources must be shared among new development, refactoring, bug-fixing and CE. The quality of past releases, because

it determines the level of CE and bug-fixing, thus has a significant influence on the resources available for the development of new features in the current release.

Project Beta

Product Beta is part of a complicated telecommunication switch that is developed largely independent of the parent product but is tested, launched, and sold together with that parent.

While Beta includes both software and hardware elements, the majority of its over-80-person R&D resources focus on developing the software. Beta's R&D team is spread across multiple locations with four main sites, two of which overlap with Alpha.

The Beta organization follows the proscribed Sigma practice closely. Its development practice includes early review and testing, root-cause analysis, and extensive automation of testing. The specificity of requirement documentation along with comprehensive review in the coding steps enable the avoidance or early discovery of problems. Moreover, early testing of code inside the development organization is prevalent in Beta, resulting in a small number of bugs once the product is sent to the quality assurance organization for integration and performance testing. Field problems are sufficiently infrequent that when one does appear, the development organization can often do a root-cause analysis to determine its sources, thus identifying and eliminating a potential failure mode in the underlying process.

The Beta product is recognized throughout the organization for its high quality and consistent delivery. While in other projects it is not unusual for thirty percent of the development resources to be engaged in current engineering activities, in Beta this number has always remained under ten percent. Beta has also delivered most of its releases on time. In two cases Beta had to delay

a release to remain in synch with its parent product, which was delayed. Beta's quality and size are reported in **Figure 1**.

<Insert figure 1 about here>

Alpha

Product Alpha is a Customer Relationship Management (CRM) software system used to manage call centers. Alpha was initially produced by a small start-up firm that was then acquired by Sigma. Alpha has been part of Sigma for 8 years and the team working on it has fluctuated around 120 full-time employees who work in five main locations. We focused our data-gathering effort on the 4 recent years of its history. Alpha's initial releases met with significant success, including leading its market, and, at the time, Sigma considered Alpha to be an important and promising product. However, during the last two years we studied, due to long delays in delivery and low quality, Alpha ceded its leadership position. The decline has been so significant that its long-term viability is in question.

In contrast to Beta, Alpha's use of best practice has eroded with time. While the official development process calls for comprehensive requirement review, development of detailed requirements, code review sessions in which individuals review each others' work, and multiple testing and documentation steps throughout the development, Alpha's developers often skip these steps. Experienced developers report that the level of requirement review is low and a growing fraction of the development work starts based on aggregate requirements (rather than the detailed requirement specified by the Sigma development process). This practice creates the danger of missing potential interactions among modules of code, errors in the interfaces between different modules, and complications in integration testing.

Code review and unit testing have also suffered, leading to a growing chance of problems surfacing later in the testing process and unexpected surprises. Documentation of code is now far from complete and developers have a hard time building on old modules, which are often poorly documented. While reviewing and testing code early in the development process is initially, research suggests that these investments yield a favorable return due to the avoided rework that would otherwise be necessary when problems are identified through quality assurance or by customers (Jones, 2000). Alpha's current practice no longer includes many of these basic activities and consequently its performance has deteriorated on several important dimensions, as reflected in the comments of an experienced developer:

“What I have seen is years and years of degradation – getting worse and worse over time.”

Alpha's performance deteriorated in each of the three categories in which software development is typically measured: schedule, quality, and productivity. Timeliness measures how well a project has followed its schedule and has a strong impact on profitability (Ethiraj, et al., 2005). Measured either in absolute terms or as a percent of the project's original schedule, Alpha's timeliness worsened or did not improve over the six releases we studied (**Figure 2-a**). Quality, typically defined by customer's experience of the software, can be measured by the number of serious bugs reported per customer per unit of time in the first few months of installation of the software. As shown in **Figure 2-b** and **2-c**, the number of bugs in Alpha grew on average over the period of study. Finally, the productivity of the team can be measured by the number of lines of code developed per individual involved with the R&D process. As shown in **Figure 2-d**, Alpha's productivity declined significantly over the second half of the study period.

<Insert Figure 2 about here>

Research Questions

The histories of Alpha and Beta pose two challenges to the current understanding of organizational capabilities. First, although Alpha was deemed strategically significant and therefore received the continued support of Sigma's leadership, its performance demonstrably declined over the period of study. Given the lack of a significant technological or economic disruption during this time, existing theory offers little to explain why. To the contrary, the current view holds that organizations are fundamentally adaptive (Nelson, and Winter, 1982), resulting in ongoing learning and improvement (Argote, and Epple, 1990). While routines are recognized to be susceptible to convergence to local maxima and therefore often impede an organization's response to technological or environmental shifts (Levitt, and March, 1988; Leonard-Barton, 1992; Levinthal, 1997), and absolute decline in the effectiveness of routines in a stable environment is not predicted by current theory.

Second, although Alpha and Beta shared many attributes including, size, resources and organizational and incentive structures (several of the people involved on the two products even worked in neighboring office), their results differed. A researcher in Sigma familiar with both products summarized the state of the two development groups by saying:

“Projects [Alpha] and [Beta] are significantly different in their development processes and their results. [Alpha] is a project that has been in continual turmoil, producing much code, often late and of poor quality and not selling well in the market. It must devote considerable resources to fixing field errors. [Beta] ... is almost always on time and produces high quality code with few field errors that sells well in the market”

Due to random changes in the customer demands, personnel turnover, and other exogenous factors, transient differences in efficiency and practices that constitute development capability are expected across similar organizations. One would not expect these differences to be permanent, however, as neighboring organizations learn from the experience of others and

converge to similar routines and performance. In contrast to the current view, Alpha and Beta, despite similar organizational and technical conditions, displayed diverging trajectories of both performance and practice. In the next section, we use this empirical anomaly to theorize about dynamics of capability evolution, specifically addressing the question of how differential performance can emerge in product development capability and why it persists. Based on these two cases we use an inductive approach to build a formal model of the dynamics that govern the evolution of development capability and explore how these dynamics can lead to heterogeneous performance across different organizational units.

Model and Analysis

We develop our answer to these questions in two pieces. First we propose and analyze a model to explain the declining performance of the Alpha project. Building on a simple representation of the “physics” of developing software, we show how the intendedly rational decisions of participants within that process combined to create a phenomenon that we call the *adaptation trap*. The adaptation trap details how the well-intentioned efforts of managers to optimize the balance between the demand for new features and the resources to develop those features can inadvertently trap the project in a vicious cycle of declining productivity and increasing demands. Second, we interpret the Beta project’s practice within the context of the adaptation trap and show how Beta’s managers successfully avoided trap that proved so costly for Alpha.

The Rise and Fall of Alpha: The Adaptation Trap

We develop the concept of the adaptation trap in three steps. First, we show how the basic physical structure of any work process when coupled with the effects of work pressure on productivity creates a system with a tipping threshold, a level of activity that once exceeded

causes a sustained decline in performance. Although it is derived from very different assumptions, this model is structurally similar to that offered in Rudolph and Repping (2002), thus we present it briefly and refer the reader to the original paper for more details. Second, we extend the model to capture how the dynamics change when the work process in question creates an installed base of products that can require additional attention (e.g., bugs that need to be fixed). These additions both add to the realism of the model and help explain the persistence of differing performance between Alpha and Beta.

Work Pressure and Tipping. **Figure 3** shows the stock and flow structure of our initial formulation. The stocks (boxes) represent the integration of different flows (thick valves and arrows) and help distinguish between a particular action or activity (represented by a flow) and its accumulated impact (represented by one or more stocks). Stock flow networks constitute the foundation of any dynamic system (Sterman, 2000) and are central to characterizing its behavior. The first construct in our model, *Demand for New Features*¹, represents the flow of requests for additions to the system. In both Alpha and Beta, new features were proposed by sales, services, and marketing groups, and further augmented by the R&D managers. While the demand for new features is largely exogenous to the R&D department, product managers did have some flexibility in prioritizing among potential features (see **Figure 3**).

The flow of newly proposed features accumulates in the stock of *Features Under Development*. Feature requests remain in this stock until they are developed, tested, and released in the market to become part of the *Features in the Current Release* of the software. The flow *Feature Release* captures the speed at which the PD organization develops and releases new features and therefore how often new releases are launched.

¹ The names of variables that appear in the model diagrams are *italicized* when they are first introduced.

<Insert Figure 3 about here>

Capability manifests in this setting as the ability to transform requests into defect-free features that appear in the next release of the product. Companies that sustain high rates of feature release for a given resource base will have a competitive advantage through some combination of more features, faster release times, higher product quality, and lower cost.

It is difficult to attribute Alpha's decline to exogenous factors both because there were few such changes over the study period and, during that time, Beta, which faced similar constraints, flourished. Thus, to understand why the performance of the two projects diverged over time, we focus on the management of the project. Our data suggest that a key difference between the two projects was the number of pending requests for new features. Beta tended to face fewer such requests and was not willing to adjust its approach to development when the demand for new features exceeded its ability to deliver. Alpha, in contrast, tended to face a larger stock of new requests and, importantly, often attempted to respond to those requests. An Alpha manager elaborates on this difference:

“[Beta Group] looks at what are the critical features, what are the extra features, and they pull features out [if unrealistic] We are not good at that. “

To model the differing approaches to managing the work-resource balance we begin with the notion of the *Resource Ratio (RR)*, defined as the ratio of resources required to finish the development of the pending Features Under Development by a scheduled date to the available resources. The Resource Ratio is a function of three related decisions: how many features are under development, how much time is available to develop those features, and how many development resources are available to do the job. The Resource Ratio is a measure of the *pressure* on the development team to deliver their work more rapidly and is directly (and

negatively) related to organizational slack (Bourgeois, 1981). When the ratio is low, the team can meet its targets relatively easily; when the ratio is high achieving targets is difficult. Pressure can have a significant effect on the speed at which people work (Mandler, 1984;Fisher, 1986;Gary, 2005). Abdel Hamid and Madnick (1991) estimate that in software engineering people can work as much as 40% faster when under pressure. Such flexibility is valuable as it enables organizations to accommodate unplanned variations in workload. Both Alpha and Beta adjusted work pressure to keep projects on schedule, although as we will discuss later, Alpha resorted to this approach far more often.

To capture this positive effect of work pressure in **Figure 4** we show the first feedback process in our model, the *Get the Work Done* loop. As the resource ratio grows, pressure increases, leading to a higher rate of *Feature Release*, which, in turn, reduces the stock of *Features Under Development* and therefore alleviates the pressure. All else being equal, this process, a balancing feedback (denoted by the B in the loop's center), works to offset any imbalance between resources and requests, thus enabling development teams to accommodate variations in their workload without constantly adjusting the size of the team or changing the delivery schedule.

Work pressure also has a dark side: as pressure increases and developers respond by working faster, they start to make more errors. The increased speed often comes through shortcuts and less attention to detail, leading to inadequate requirement development, little documentation of the code, lack of code review, and poor unit testing (see (MacCormack, et al., 2003) for a quantitative analysis of the effects of different practices on quality and productivity). Moreover, stress induced by sustained pressure can reduce individual productivity and increase absenteeism

and turnover, further limiting output. A sustained, high, resource ratio can therefore be counterproductive. In the words of a technical manager in Alpha:

“There is ... a lot of pressure put on people to say how can you do this as fast as you can, and people fall back on the ways that they used to do it, so instead of going through all this process stuff [they would say] ‘I know what to code, I can just code it up.’”

We capture the negative effects of work pressure by adding a second feedback to our formulation. As resource ratio (i.e. work pressure) grows, the error rate (broadly defined to include anything that limits net productivity) increases, thereby slowing the net completion rate. As the completion rate slows, the stock of pending features grows (assuming a constant inflow of feature requests) thus feeding back to create even more work pressure. Unlike the “Get the Work Done” loop, this new process, the “Make More Mistakes” loop, is a reinforcing feedback that amplifies any given disturbance. When working in the downward direction— a growing error rate, declining net output, and mounting resource ratio—the “Make More Mistakes” loop acts as a vicious cycle.

<Insert figure 4 about here>

Analyzing the dynamics created by the interplay of these two processes requires specifying the net impact of pressure (Resource Ratio) on the rate of Feature Release. Consistent with psychological research on Yerkes-Dodson law (Mandler, 1984), and previous research in product development (Taylor, and Ford, 2006) and software development (Abdelhamid, and Madnick, 1991), our discussions with study participants pointed to an inverse U-shape relationship between the resource ratio and net output. At modest levels, resource ratio (pressure) tends to improve net productivity as the positive effect resulting from working faster offsets any increase in the error rate. However, the positive effects of resource ratio decline with increased pressure,

while the negative effects tend to grow, thus causing the net effect to eventually become negative.

As we mentioned above, while it is derived from a very different set of assumptions, our initial formulation is structurally similar to that described in Rudolph and Repenning (2002), so we only summarize the basic results before moving to our extended formulation. Figure 5-a shows a characterization of the system's dynamics. We begin with the inverted U-shape curve that relates the current resource ratio to the feature-release rate and then characterize the system's dynamics by showing both the system's equilibria and its behavior around those equilibria.

<Insert Figure 5 about here>

The system has two equilibria, each of which occurs at the intersection of the rate of feature demand and the inverted-U. These equilibria represent the points at which the resource ratio induces just enough work effort to meet the current level of feature demand. Once the system reaches either of these points, it will be in steady state (i.e. equilibrium) with the rate of feature release equaling the rate of feature demand. Although both points represent equilibria the dynamics around them are very different. The equilibrium that appears on the upwardly-sloping portion of the inverted-U is *stable*, meaning that small perturbations from it are offset by the system's dynamics. For example, a small increase in the resource ratio (perhaps due to a temporary increase in the stock of features pending) would cause the feature completion rate to increase (manifesting as a temporary excursion off the stable equilibrium and up the inverted U) above the feature demand rate and thus reduce the stock of features pending and bring the system back into balance. Formally, in this region the balancing Getting Work Done loop dominates the system's dynamics and functions exactly as managers desire, constantly working to bring the level of outstanding work and the available resources into balance.

In contrast to the equilibrium in the upwardly sloping portion, the equilibrium that occurs in the downwardly sloping portion of the inverted U is *unstable*, implying that any deviation off that equilibrium, rather than being offset, will be reinforced. Formally, this equilibrium is unstable because it occurs in a region where the system is dominated by reinforcing feedback. As indicated by the arrows on the inverted-U, the system is unlikely to ever settle at this point since even the smallest perturbation will push the system on to a new trajectory, away from the unstable equilibrium. Due to these dynamics, the unstable equilibrium is often called a *tipping point*, capturing the idea that once an unstable equilibrium is crossed, the system's dynamics can change significantly.

Consistent with this intuition, the tipping point plays a key role in determining the dynamics of the system we study. **Figure 6** shows an experiment in which the Feature Demand rate has been temporarily increased. In the first experiment (**Figure 6-a**) the shock is not sufficient to push the system over the tipping threshold (both experiments start with the system in steady state, operating at the stable equilibrium) and consequently, the system is able to absorb the shock and return to its initial performance level (as measured by the feature release rate). In the second experiment (**Figure 6-b**), in contrast, the increase is large enough to push the system over its tipping threshold. Here, rather than returning to the initial equilibrium, the system gets stuck in a vicious cycle of growing feature demand, mounting resource ratio and a declining rate of feature release. As the experiment indicates, when the system crosses its tipping threshold, performance declines significantly and does not recover.

<Insert Figure 6 about here>

Our characterization highlights three key features of the system's dynamics. First, due to the existence of the tipping threshold, the system's response to temporary shocks is non-linear. Any shock that does not cause the system to cross the tipping threshold is offset, while shocks that do exceed that limit cause a permanent decline in performance. Second, the size of the shock required to push the system over the tipping boundary is determined by the distance between the stable equilibrium and the tipping point. When the distance is large, the system can absorb a substantial shock; when it is small, a modest perturbation can push performance over the edge. Third, the distance between the two equilibria is a function of the steady state rate of feature demand. As shown in Figure 5-b, as the rate of feature demand rises, the two equilibria move up the inverted-U and the distance between them shrinks. If feature demand continues to rise, eventually both equilibria disappear (the horizontal line lies strictly above the inverted-U), and the system displays only one trajectory, a downward spiral into low performance. Thus, managers in this system face a clear trade-off: As they increase the steady state-level of feature demand, performance continues to improve (as long as they stay within the inverted-U). However, as feature demand rises, the size of the shock required to push the system over its tipping threshold declines.

Our data clearly indicate that Alpha's performance declined over the period of study and the evidence also suggests that they were increasingly trapped in a vicious cycle of accumulating outstanding work, mounting work pressure and declining performance. Moreover, the interviews strongly suggest that the declining performance resulted from changes in the way the work was executed. Several developers noted that, due to schedule pressure, the Alpha team increasingly abandoned the elements of good software practice such as documentation in favor of getting the work done more quickly. And, as the graphs in the previous section indicate, they suffered the

consequences. While it started well, Alpha's quality declined significantly and later releases were plagued with bugs requiring additional resources for remediation.

One might have expected that Alpha team would have recognized that they were at the risk of tipping into a vicious cycle and have reduced the feature request accordingly. There is, however, no evidence that such adjustments were ever made. To the contrary, as the tipping model would suggest, their work load only increased over the period of study. To explain why, we turn to the concept of the adaptation trap.

Bug Fixing and Current Engineering. So far we have assumed that errors are discovered within R&D department, and thus affect the system's dynamics only by reducing net productivity. The real situation in software engineering (and many other development activities) is, however, more complicated. At Sigma some errors escaped the quality assurance process altogether and remained in product that was shipped to customers. When discovered by customers, these errors would be referred back to the R&D department. If the problem interfered with an important customer need, R&D would assign engineers to fix the problem, usually through quick, customized patches at the customer site. These activities, called *current engineering* (CE) can take a significant share of R&D resources: Alpha allocated 30-50 percent of development resources to CE work; Beta, in contrast, kept this number below 10 percent.

Defects, the current engineering they require, and the delays inherent in identifying those resource needs introduce two important wrinkles to our analysis. First, CE requires additional resources. Second, and more importantly, those resource requirements can only be known with a significant delay. Customers must buy, install and use the new release before CE resource needs

are revealed. In this section we show how the addition of these two features makes the managers' job more difficult and the system's descent into firefighting more likely.

<Insert Figure 7 about here>

Figure 7 shows the structure of our expanded model. The flow of Feature Release now accumulates in a stock of *Features in the Current Release*, representing the installed base of code released to Sigma customers. A new flow, *Defect Introduction*, now runs in parallel to Feature Release, and captures the stream of defects in new features that are not identified by in-house testing. The defect flow accumulates in the stock of *Defects in the Current Release*, representing the stock of defects in software currently being used by customers.

Defects in the field, as they are eventually discovered, create additional demand for PD resources via current engineering (*Resource Demand for Current Engineering*). As we show in the diagram, the demand for CE increases the Resource Ratio, potentially creating more pressure and a higher Error Rate. These new links add an additional positive feedback to the model, the reinforcing Current Engineering loop. As the stock of defects grows and customers request additional Current Engineering, developers face even more demands on their time and must work faster to meet their delivery targets. As we discussed in the previous section, the increased work speed can result in an even greater error rate, causing the stock of defects in the field to grow further. Note that the strength (or gain) of this loop is potentially high as current engineering is often extremely costly relative to fixing problems before they are released (Boehm, 1981). Whereas an error prevented in the design phase can prevent a large number of customer complaints, current engineering activities are often focused on resolving the problem of a single customer. Detailed model equations are available in online appendix.

The addition of current engineering and the delays associated with perceiving its full impact increases the challenge facing those charged with managing resources for two reasons. First, CE requires additional resources and thus reduces the capacity of the organization to develop new features. This shifts the inverse U-shape graph in Figure 5 down. While the existence of the shift is clear, accounting for its magnitude creates a second challenge. Consider the decision to increase the rate of feature demand: Due to the delay in perceiving the increased CE demand, in the short run the system behaves as though no additional errors (from additional features) make it to customer site. Therefore the delay increases the apparent capacity of the system. In fact, as the fraction of errors that go undiscovered grows, so to does the apparent capacity of the system to do additional new work (since resources do not appear to be needed for rework or additional current engineering) (see Figure 8). In other words, in the short-run the system will behave as though it has more capacity than is actually available.

<Insert Figure 8 about here>

This feature of the system—that its short term behavior differs significantly from the long run performance—makes learning difficult. Recall from above that the distance between the stable equilibrium and the tipping threshold is determined by the level of resource utilization.

Increasing the rate of feature demand has the effect of pushing both the stable equilibrium and the tipping point “up” the inverted U, reducing the distance between them and making the system more susceptible to firefighting. Thus, the task facing managers of such systems is to set the capacity/resource balance in a way that appropriately trades-off the higher level of performance that comes with a higher level feature demand against the decline in the system’s robustness to temporary shocks.

Such a decision is not particularly amenable to formal analysis. The organization's true capacity is difficult to estimate as is the workload associated with a given set of features. Managers thus very likely engage in local search, progressively increasing the workload and monitoring the performance of the system to determine whether the increases yield a higher rate of feature release. Such local search processes are believed to be a principal mechanisms through which organizational routines adapt and create better performance (Nelson, and Winter, 1982).

Tightening the schedule to get the maximum feasible performance was a common strategy in Sigma. A program manager comments:

“The thinking that [a higher manager] had, and we think we agree with this, is that if your target is X and everybody is marching for that target, you fill the time for that target. If you are ahead of the game, people are probably going to relax a little bit: because we are ahead of the game, we don't need to work so hard, we can take a longer lunch, we can do whatever; and guess what: Murphy's law says you now won't make that target. So [...] you shoot for X minus something, you drive for X minus something so that you build in on Murphy's law to a point.”

The assumption of localized search coupled with a system that displays different short and long run dynamics leads to the notion of the *Adaptation Trap*. As shown in Figure 8, imagine a manager attempting to set the appropriate balance of work and resources by incrementally increasing the rate of feature demand. In the short-run, the system will behave as described by the upper inverted-U (which we label the apparent system). As we show on the graph, such an adaptation path will lead her to choose a level of loading somewhere near the peak of the apparent curve. However, if the true, long run, behavior of the system is described by the lower curve, then the adaptation process will have converged to a rate of feature demand that is beyond the sustainable capacity of the PD organization. When the additional CE work does eventually start to appear, the PD organization will find itself past its tipping threshold and on the slippery slope of firefighting. The tight schedule and high feature demand that initially seemed so

effective will cause the rate of feature release to fall below feature demand, leading to continued growth in the stock of features to be developed and, thereby, ensuring that pressure stays high and that the organization remains in the firefighting mode. Thus, in contrast to early experiments in which an exogenous shock was required to push the system over the edge, the Adaptation Trap notion suggests that such a decline can occur *endogenously*, resulting from the well-intentioned efforts of managers to optimize their system: due to delays inherent in getting feedback, managers do not receive complete information about the results of their experiments, leading them to systematically overload their development system.

The existence of the adaptation traps depends on managers not fully accounting for the delays in the system. Were they to be patient in assessing the full consequence of each incremental increase, they would be unlikely to overshoot. There is however significant experimental and field-based evidence to suggest that this is unlikely. On the experimental side, subjects have been repeatedly shown to both underestimate time delays and fail to properly account for their impact on a system's dynamics (e.g. Serman, 1989; Diehl, and Serman, 1995). Similarly, a growing collection of field studies document a similar phenomenon (Oliva, and Serman, 2001; Repenning, and Serman, 2002).

At Sigma, the calculation of expected resource needs for CE was even less complicated: development teams were expected to dedicate no more than 10% of their workload to CE activities. This was the accepted norm across different PD groups and was a key part of the resource allocation process for all projects. We observed that its application to planning process was hard to challenge even in the presence of contrary data. For example the same individuals who described the real current engineering load to be over 30% in the case of Alpha, also

explained that they planned about 10% of resources for current engineering. Such simple rules are a common and often efficient response in complex and uncertain dynamic systems. Note, however, that the application of such a rule effectively prevents Alpha from ever escaping the firefighting trap. Because they underestimate the resources needed for CE, sometimes by as much as a factor of five, they take on too much new work, thus insuring that they remain on the wrong side of the tipping threshold.

The strength of the adaptation trap depends on three features of the system, the delay in perceiving the required level of CE, the speed with which managers attempt to optimize resource loading, and the relative productivity of CE versus normal development work. A longer delay in perceiving the demand for CE resources results in a longer window over which adaptive processes work in the absence of the CE feedback, thus increasing the risk of tipping. Similarly, faster adaptation implies that managers will use less information in determining whether a given change produced a gain in productivity. As managers shrink their evaluation horizon, they account for a smaller portion of the CE effect. Finally, as the productivity of CE falls (relative to normal development work), the trap becomes more likely since each defect that makes it to the field becomes more costly in terms of the foregone development resources. With release cycles (and thus CE feedback delays) of up to one year, resource pressures that pushed for full utilization (thus fast adaptation), and costly customized CE, Alpha faced something of a “perfect storm” in managing its workload.

Worse, although we have so far discussed the adaptation trap as though the project was starting from steady state, this was not the case for Alpha, a state which further exacerbated the challenge posed by the adaptation trap. In developing the first release of a product there is, by definition,

no CE demand. In fact CE demand remains low as long as there are few customers for the product. Thus in its early days, the Alpha team could spend most of its resources on new development activities, leading to a high estimate of the amount of new work it could accomplish. As Alpha moved to subsequent releases and sales grew, the demand for current engineering began to grow. In the meantime, Alpha was well received by the market and was installed by several high-profile customers, leading Alpha's management to request additional features focused on maintaining Alpha's reputation as an innovative product. Later, as CE requests arrived, the team was careful to respond to them quickly to keep a good relationship with its highly visible customers. The Alpha team thus soon found itself significantly overloaded and was forced to depart from Sigma's normal development practice in an effort to meet the mounting demand for new features and CE.

Note also how abandoning standard development practice coupled with the delays in discovering CE further constrained Alpha's adaptation process: as developers took short cuts, apparent productivity increased because many of the consequent errors would only be discovered when the product reached the customer site, often several months after that work was completed. Thus the changing work habits of the development team made it even more difficult for managers to assess the team's true capacity. As one developer reported:

“Nobody really sees what we do. You get rewarded when your name is next to a dollar sign Nobody knows what actually the quality is and what actually we do.”

The result of these dynamics was a product that was in part a victim of its own success. Alpha's successful market reception coupled with the delays in assessing the true quality of its code led managers to overestimate the team's capacity for new work. The subsequent workload

eventually pushed Alpha over its tipping threshold creating the situation highlighted in the previous section, declining quality and productivity.

The Differing Trajectories of Alpha and Beta

Our analysis suggest that in a well-intentioned effort to maintain its market leadership, the Alpha team failed to accurately assess its true capacity and found itself trapped in a vicious cycle of mounting delivery pressure, declining performance, and an increasingly unwieldy and confusing product architecture. Given our earlier claim concerning their similarity, our analysis immediately raises of the question of why did the two projects produce such differing levels of performance? Whereas Alpha suffered a serious decline, Beta continues to be one of the most successful projects at Sigma. Despite their many similarities, Alpha and Beta differed on a few key dimensions, and as we outline in this subsection, those differences, when coupled with the basic structure that we have described so far, caused the two projects to experience very different performance trajectories.

Structural Factors. Our account of the differing performance trajectories of the two projects begins with a key difference between them. While Alpha was a stand-alone product, Beta's delivery was gated by its parent product, a telecommunication switch. Thus, whereas Alpha was directly subjected to delivery pressure from the market, as long as Beta was making more progress than its parent product, the pressure it felt was moderated. Two times in Beta's history, the parent product was delayed even though the Beta team was on track to hit the original schedule. Delays in the delivery of the parent product provided the Beta team with time to reflect, improve their processes, and prepare for the next release. Each of these delays pulled

Beta farther below its tipping threshold, allowing the development team to increase the effectiveness of their work. One of the Beta managers reflects on this experience:

“How we got here? The answer is in investing in [the] future. And from the [Beta’s] perspective we got lucky. If I go back to the first release [of Beta],... they [the parent product] were running into many more schedule problems than we were ... this allowed us to go back and actually do a lot of things that we would have preferred to have done at the beginning but we never got the time to do... So we benefited by that additional time to really add a lot of foundational components.”

Note how this advantage, once it tipped towards Beta, could be self-reinforcing; the more the parent product struggled, perhaps due to being in a situation similar to Alpha’s, the more time Beta had to focus on high quality work, thereby creating even more space for itself. Alpha, in contrast, benefitted from no such protection. It continually faced market pressure and, as its market position slipped due to quality problems, there was additional pressure to add new features to compensate, creating yet another vicious cycle.

Initial Conditions. Alpha and Beta also had different origins. While Beta was developed internally by a team with significant experience, Alpha was born in a start-up organization, subsequently acquired by Sigma, largely staffed by younger, less experienced managers and developers. Their differing antecedents influenced how the two projects were managed. The Beta team had significant experience managing a product over its full life-cycle. Many team members reported having past experience with the costs of skimping on the development disciplines and therefore placed a strong emphasis on careful adherence to good practice. One of the Beta managers recalled:

“Early in my career... when I was in an organization that did wireless... we were under tremendous market pressure to get the product out...and you just do anything you can to get anything out, not one comment was put in a lot of code, not one thing was inspected, not one thing was documented, you know, it was classic, and it never worked, and still to this date it is in the field, and it has problems... and that’s were I really cut my teeth and learned my lesson. In [Beta]... we have lots and lots of huge pressures... and we don’t skip process,

because that is the tone we set... I have had people who skip process and I have had them fired, that it is not a platitude, I insist.”

Alpha, in contrast, initially faced all of the pressures typical of a start up venture. Lacking an ongoing revenue stream and hoping to realize the benefits of first mover advantage, the team felt significant pressure to get their product to market. And, once that initial release proved successful, the demands for additional features only added to the pressure. These conditions reinforced a culture of “quick and dirty” problem solving, which although possibly helpful in the beginning, proved counter-productive later in the project’s life.

Management Policies. The differing initial conditions and external environment combined with the dynamics we have discussed to create two very different approaches to managing software development. Discussions with Beta’s managers suggested that many of them had an intuitive sense of the dynamics we discuss and thus actively worked to avoid them. Notably, Beta displayed several organizational practices targeted at maintaining quality that were not present in Alpha. For example, root cause analyses of problems reported from the field were often conducted in Beta, and the results could impact personnel review and incentives. The developers therefore knew that if they cut corners and things went wrong, they would be held responsible when bugs surfaced in the field. Similarly, the Beta organization instituted a “No-Late-Feature” rule, which was enforced by the team’s leadership. Whereas Alpha was frequently subjected to changes in requirements far into the design phase, Beta’s managers were typically able to resist the pressure to add features once design started. The rule helped avoid spikes in workload not captured in the original plan and thus played a central role in keeping Beta below its tipping threshold. Finally, in contrast to Alpha, Beta had a separate, dedicated team of developers focused on current engineering. By not allowing the increase in CE demand to impact the

development resources available to the next release this policy effectively breaks the current engineering loop discussed above.

At a first glance, it might appear that, in contrast to their colleagues in the Beta project, Alpha's managers were simply naïve and unskilled. A closer examination suggests that the situation is more complicated. Alpha's managers reported being able to see the benefits of Beta's practices (which reflected industry best practice), but they found them infeasible for their own organization. For example, root-cause analysis is very time consuming and only possible if the project faces a relatively small stream of bugs coming from the field. While Beta was able to do a full root cause analysis for many of its bugs, Alpha's bug stream was more than an order of magnitude larger, thus making comprehensive root-cause-analysis infeasible.

Attempting to separate the resources dedicated to CE creates similar challenges. CE demand is uncertain and hard to plan. When the CE demand is low, it is possible to dedicate a small team to that work and accept the inevitable idle time that occurs when no new bugs are reported. When CE work grows to over 30% of the total resource load (as it did in Alpha), it becomes prohibitively expensive. The organization simply cannot bear the costs of keeping a significant fraction of its people idle at the times that demand is low. Both examples suggest that team's approach to managing the project tended to co-evolve with the state of the project itself. Beta's stability tended to be self-reinforcing, creating more space for the investments necessary to maintain best practice. Alpha, in contrast, faced increasing pressure for short-term delivery, causing it to progressively abandon high performance practice.

The "No-Late-Feature" rule provides an example of how this co-evolution spreads beyond the internal workings of the team to the surrounding organization. Several informants reported that

one of the main reasons that the Beta team could resist the pressure of senior management and customers to add late features was the reputation and bargaining power gained from their history of high quality and on-time delivery. In absence of such credibility and bargaining power, Alpha's managers found it far more difficult to resist late feature requests, which they generally agreed were detrimental to the project's performance.

Thus, although both teams were familiar with the disciplines that constituted effective software development, as the two projects progressed, the teams' abilities to put those disciplines into practice diverged. Whereas the Beta project continued to follow best practice, Alpha found itself increasingly trapped in a vicious cycle of mounting quality problems, increasing pressure for additional productivity and declining adherence to best practice. Our study therefore suggests that luck, incentive structures, and managerial cognition can co-evolve, creating a set of self-reinforcing feedbacks capable of producing very different capability endowments across organizations. While previous research has emphasized serendipity (Denrell, et al., 2003), local peaks (Levinthal, 1997), or superior insight (Kaplan, et al., 2003) as the core explanatory factor, we find that the dynamics governing their co-evolution may provide deeper insight into the question of why heterogeneous capabilities evolve and persist. Such view puts the identification of the feedback loops governing these co-evolutionary processes at the forefront of the research: heterogeneity is most likely when reinforcing loops dominate.

DISCUSSION

Comparing the experience of Alpha and Beta thus raises a key question for students of strategy and organization: what is the relevant capability that Beta managed to maintain but Alpha did not? Consistent with previous work (e.g. Klein, and Sorra, 1996; Pfeffer, et al., 2000; Repenning,

and Sterman, 2002) our data suggest that knowledge of the specific practices did not, in and of itself, constitute a relevant capability. Alpha's managers were familiar with the activities that constituted good software practice and frequently lamented their inability to use them.

Moreover, members of Sigma's central development organization had in some cases literally written the book on good software development practice; *knowing* the right thing to do, while necessary, was not sufficient to produce high performance in every project.

Comparing Alpha and Beta suggests that the ability to translate knowledge into effective practice both turned on the state of the project in question and tended to be self-reinforcing. The relatively high quality of the Beta product afforded the development team both the space and the authority to follow best practice and thus maintain the quality of their product delivery. In contrast, Alpha's troubled performance eroded the team's power to resist the addition of features and increasingly pushed the team to abandon the disciplines associated with good practice. Such reinforcing processes appear to be at the heart of differing capability trajectories.

Identifying the positive feedbacks at the heart of capability development and erosion is, however, only part of the answer. What policies or structures dictated whether these processes worked in an upward, virtuous direction, as in Beta, or a downward, vicious direction, as in Alpha? Our search for the answer to this question returned us to a notion at the very foundation of organization theory: the ability to balance resources and demands. At the core of our account of the two differing trajectories lies one specific decision point, the amount of work undertaken given the available resources. Analyzing our model suggests that when teams try to take on too much work, they run a significant risk of tipping into a pathological regime in which it is increasingly difficult to follow good development practice. While central to early accounts of

organizational behavior (Pfeffer, and Salancik, 1978), the notion that organizations should balance resources and demands has received little attention in the recent literature. Despite the lack of scholarly attention, our analysis suggests that under conditions to be outlined momentarily, balancing resources and demands is central to allowing an operational capability to flourish. Our data suggest that when organizations take on more work than they have the resources to accomplish, participants change their behaviors in ways that cause capability to erode. And, although it has not been previously highlighted as a capability, the criticality of resource balancing in insuring the continued use of the best practices has been highlighted in several previous studies (e.g. Repenning, 2001; Taylor, and Ford, 2006).

The first contribution of our study is to highlight the ability to balance resource and demands as a *foundational capability*. The skill and ability to write quality code is also certainly a capability, but our data suggest that it is less relevant in explaining the differences in performance than the contextual and managerial factors that explain its use or lack thereof. Distinguishing between so-called zero-level capabilities—the ability to develop code—and foundational capabilities—the managerial skill and foresight to balance resources appropriately —draws attention away from specific practices and technologies, which vary widely across industries, to basic managerial and administrative functions. Redirecting attention in this manner may help strategy scholars hone in upon a more general set of principles from which to understand the sources of competitive advantage.

Highlighting the criticality of resource balancing as a foundation on which capability is developed does appear at odds with some of the current literature on high performance organizations, particularly software development. Recently, some authors have argued for the

utility of methods such as agile development in which careful planning and resource balancing are scrapped in favor of rapid iteration. In such methods, development teams simply work through as many build-test cycles as is possible prior to the release of the project and planning plays a relatively minor role. Anecdotal evidence suggests organizations following these methods can also produce above average performance (Martin, 2002; Larman, 2003).

Within the resolution for these competing views lies the second major contribution of our study. Our efforts have identified a specific, foundational capability that may eventually be shown to be a general source of competitive advantage. That said, we do not claim that resource balancing will be strategically important in all contexts. To the contrary, our formal analysis yields a clear set of conditions under which the pathological dynamics we study are likely to appear. As such, our analysis yields three sharp hypotheses concerning when resource balancing will constitute a strategically important capability.

The first of these structural conditions concerns the role of work quality in creating additional demands on the organization. The tipping dynamics we study turn heavily on the fact that low quality work fed back to create additional work and put additional pressure on the development team. When this assumption is not satisfied, the positive loops at the center of our analysis do not exist. Thus, the first hypothesis emerging from our analysis is that the value of resource balancing is positively associated with the resource requirements resulting from poor work.

In the cases of Alpha and Beta the link between poor work and subsequent resource requirements was strong; each defect could create a significant demand for both current engineering and bug fixing. Not all development activities, however, satisfy this condition. For example, in many industrial settings, development teams simply have to live with the defects in the products they

create and have few opportunities to fix them. Toys and fashion goods are examples. Although defects in the design may cost the company money in terms of sales and/or warranty expense, the designers and engineers in this context are unlikely to be distracted from their work on the next generation since the designs are almost entirely new from one year to the next..

The second condition concerns the fungibility of resources between doing work and fixing the problems arising from past efforts. The reinforcing feedbacks we study were particularly strong because it was easy for the development teams to shift between new work and fixing bugs.

When resources are less fungible, those dynamics will be muted. Thus our second hypothesis suggests that resource balancing will yield larger gains in performance when resources are fungible. Much like the first hypothesis, software development is a strong case in point. Project resources are often perfectly fungible, with development teams working on everything from developing new features to fixing issues at specific customer sites. Heavy manufacturing, such as car production, represents an intermediate point on the fungibility continuum. When problems appear in the field, engineers need to develop a fix and insure that the defect is eliminated in future production. In contrast to Alpha and Beta, however, engineers in this context rarely go fix individual automobiles. Finally, consider the case of developing a feature film or a high fashion product. In these cases, while defects originating from poor development will be expensive, such problems are not likely to draw on the resources that created them. Filmmakers and designers do not redo their products when they are poorly received by the market; that burden falls to their respective marketing departments.

The final condition concerns the delay between performing work and perceiving its true quality. In our analysis, this delay proves central in explaining the drift towards low performance.

Because managers cannot easily anticipate future resource requirements resulting from defective work, they are prone to taking on too much work, thereby getting stuck in the adaptation trap. Thus the final hypothesis emerging from our analysis is that the value of resource balancing will rise as the time required to perceive resource requirements increase.

Here Sigma probably represents an intermediate case. Development managers at Sigma were able to see the quality of their teams' work over the course of six to twelve months. The delays are much longer in other industries. In the auto industry, for example, it can take several years to develop a new vehicle and then at least one more year before the true quality of that vehicle has been revealed (and some problems can take far longer). On the other side of the continuum, website developers (frequent proponents of agile development methods) can often get feedback on their work in a matter of days and easily make adjustments. Our analysis suggests that resource balancing is a key capability potentially explaining firm heterogeneity in the auto industry but of little explanatory value in developing websites.

Our analysis thus yields a contingent view of the value of resource balancing and in doing so provides a potential resolution to an ongoing debate in the management literature. For firms whose products tend to have long lifetimes and designs that build upon previous releases, resource balancing will be a key source of competitive advantage. In contrast, for firms whose products are less cumulative, subjected to faster learning, and whose defects do not generate additional work, resource planning will be less relevant as a core capability. In fact, for these products, resource balancing may actually impede performance as it requires formal processes that reduce flexibility and the ability to adapt to changing market conditions.

Our work has a broader implication for managers and theorists. While early students of organizational behavior recognized that resource balancing was a key function of the executive, it appears that the scholarly and practitioner communities have subsequently lost sight of this function. Three forces appear to be responsible. First, as argued by Barley and Kunda (Barley, and Kunda, 2001), organizational scholarship has grown more distant from the nuts and bolts of getting work done and few activities appear more routine than matching work to resources. A second and related trend is the rise of structural functionalism and economic-based explanations (Perrow, 1986; Pfeffer, 1993). The dynamics we identify are fundamentally pathological and would be extremely difficult (although probably not impossible) to explain with a rational or functional approach. Thus, it appears, that many have assumed that resources must be roughly in balance and not worthy of further attention. There is growing evidence however that many organizations take on far more work than they are capable of doing.

Third, this blind spot does not appear to be unique to scholars. To the contrary, management practitioners (at least those in western companies) appear to grow ever more obsessed with individual accountability and many view overload, often phrased as a “stretch objective”, as a key tool in insuring individual performance. As in Sigma, many managers now believe that the worst thing that can happen if they put too much work into the system is that their people will be fully utilized.

Ironically, this view has been thoroughly discredited in the world of discrete manufacturing. Perhaps the most significant lesson of the so-called lean revolution is that attempting to run a factory above its capacity is a sure strategy for maintaining low performance. But, while this view has been widely accepted on the factory floor, with a few exceptions, it has yet to permeate

other types of work. And, there is a credible argument that the costs of this misconception are significant. A crushing workload and time pressure have been implicated in essentially every major industrial accident in the last twenty years including the two space shuttle disasters and the explosion at the BP's Texas City refinery (MacKenzie, et al., 2007). Similarly, the growing collection of studies focused on medical error reach a similar conclusion, too many patients being served by too few doctors and nurses is a recipe for disaster. Whether these arguments will survive empirical test is an open question. But, the existing data suggests they provide both a fruitful area for new work and, more importantly, an opportunity for research to make a significant contribution to practice.

References:

- Abdelhamid, T., and S.E. Madnick
1991 Software project dynamics: An integrated approach. Englewood Cliffs, NJ: Prentice-Hall.
- Argote, L., and D. Epple
1990 "Learning-Curves in Manufacturing." *Science*, 247: 920-924.
- Barley, S., and G Kunda
2001 "Bringing Work Back In." *Organization Science*, 12: 76-95.
- Barney, J.
1991 "Firm Resources and Sustained Competitive Advantage." *Journal of Management*, 17: 99-120.
- Boehm, Barry W.
1981 Software engineering economics. Englewood Cliffs, N.J.: Prentice-Hall.
- Bourgeois, L.J.
1981 "On the measurement of organizational slack." *Academy of Management Review*, 6: 29-39.
- Davis, J. P., et al.
2007 "Developing theory through simulation methods." *Academy of Management Review*, 32: 480-499.
- Denrell, J., et al.
2003 "The economics of strategic opportunity." *Strategic Management Journal*, 24: 977-990.
- Diehl, Ernst, and John Sterman
1995 "Effects of Feedback Complexity on Dynamic Decision Making." *Organizational Behavior and Human Decision Processes*, 62: 198-215.
- Dierickx, Ingemar, and Karel Cool
1989 "Asset Stock Accumulation and Sustainability of Competitive Advantage." *Management Science*, 35: 1504-1511.
- Eisenhardt, K. M., and J. A. Martin
2000 "Dynamic capabilities: What are they?" *Strategic Management Journal*, 21: 1105-1121.
- Eisenhardt, Kathleen M.
1989 "Building theories from case study research." *Academy of Management Review*, 14: 532-550.
- Ethiraj, S. K., et al.
2005 "Where do capabilities come from and how do they matter? A study in the software services industry." *Strategic Management Journal*, 26: 25-45.

- Fisher, S
1986 *Stress and Strategy*. London: Lawrence Erlbaum.
- Forrester, Jay W.
1961 *Industrial Dynamics*, 1 ed. Cambridge: The M.I.T. Press.
- Gary, M. S.
2005 "Implementation strategy and performance outcomes in related diversification." *Strategic Management Journal*, 26: 643-664.
- Gavetti, G.
2005 "Cognition and hierarchy: Rethinking the microfoundations of capabilities' development." *Organization Science*, 16: 599-617.
- Glaser, Barney G., and Anselm L. Strauss
1967 *The discovery of grounded theory; strategies for qualitative research*. Chicago,: Aldine Pub. Co.
- Helfat, C. E.
2000 "Guest editor's introduction to the special issue: The evolution of firm capabilities." *Strategic Management Journal*, 21: 955-959.
- Helfat, C. E., and M. A. Peteraf
2003 "The dynamic resource-based view: Capability lifecycles." *Strategic Management Journal*, 24: 997-1010.
- Henderson, R., and I. Cockburn
1994 "Measuring Competence - Exploring Firm Effects in Pharmaceutical Research." *Strategic Management Journal*, 15: 63-84.
- Henderson, R. M., and K. B. Clark
1990 "Architectural Innovation - the Reconfiguration of Existing Product Technologies and the Failure of Established Firms." *Administrative Science Quarterly*, 35: 9-30.
- Jones, C
2000 *Software assessments, benchmarks, and best practices*. Boston, MA: Addison-Wesley.
- Kaplan, S., and R. Henderson
2005 "Inertia and incentives: Bridging organizational economics and organizational theory." *Organization Science*, 16: 509-521.
- Kaplan, S., et al.
2003 "Discontinuities and senior management: assessing the role of recognition in pharmaceutical firm response to biotechnology." *Industrial and Corporate Change*, 12: 203-233.

- Keil, T.
2004 "Building external corporate venturing capability." *Journal of Management Studies*, 41: 799-825.
- Klein, K. J., and J. S. Sorra
1996 "The challenge of innovation implementation." *Academy of Management Review*, 21: 1055-1080.
- Larman, Craig
2003 *Agile and Iterative Development: A Manager's Guide*: Addison-Wesley Professional.
- Lavie, D.
2006 "Capability reconfiguration: An analysis of incumbent responses to technological change." *Academy of Management Review*, 31: 153-174.
- Leonard-Barton, D.
1992 "Core Capabilities and Core Rigidities - a Paradox in Managing New Product Development." *Strategic Management Journal*, 13: 111-125.
- Levinthal, D. A.
1997 "Adaptation on rugged landscapes." *Management Science*, 43: 934-950.
- Levitt, B., and J. G. March
1988 "Organizational Learning." *Annual Review of Sociology*, 14: 319-340.
- MacCormack, A., et al.
2003 "Trade-offs between productivity and quality in selecting software development practices." *Ieee Software*, 20: 78-+.
- MacKenzie, C., et al.
2007 "Human Factors Analysis of the BP Texas City Refinery Explosion." *Human Factors and Ergonomics Society Annual Meeting Proceedings, Safety*, 5: 1444-1448.
- Mandler, George
1984 *Mind and body : psychology of emotion and stress*, 1st ed. New York: W.W. Norton.
- Martin, Robert C.
2002 *Agile Software Development, Principles, Patterns, and Practices*: Prentice Hall.
- Mauri, A. J., and M. P. Michaels
1998 "Firm and industry effects within strategic management: An empirical examination." *Strategic Management Journal*, 19: 211-219.
- Nelson, Richard R., and Sidney G. Winter
1982 *An evolutionary theory of economic change*. Cambridge, Mass.: Belknap Press of Harvard University Press.

Oliva, R., and J. D. Sterman

2001 "Cutting corners and working overtime: Quality erosion in the service industry." *Management Science*, 47: 894-914.

Perrow, Charles

1986 *Complex organizations : a critical essay*, 3rd ed. New York: Random House.

Peteraf, M. A.

1993 "The Cornerstones of Competitive Advantage - a Resource-Based View." *Strategic Management Journal*, 14: 179-191.

Pfeffer, J

1993 "Barriers to the advance of organizational science: Paradigm development as a dependent variable." *Academy of Management Review*, 18: 599-620.

Pfeffer, Jeffrey, and Gerald R. Salancik

1978 *The external control of organizations : a resource dependence perspective*. New York: Harper & Row.

Pfeffer, Jeffrey, et al.

2000 "The knowing-doing gap how smart companies turn knowledge into action." xv, 314 p. Boston, Mass.: Harvard Business School Press.

Porter, Michael E.

1998 *Competitive strategy : techniques for analyzing industries and competitors : with a new introduction*, 1st Free Press ed. New York: Free Press.

Priem, R. L., and J. E. Butler

2001 "Is the resource-based "view" a useful perspective for strategic management research?" *Academy of Management Review*, 26: 22-40.

Raff, D. M. G.

2000 "Superstores end the evolution of firm capabilities in American bookselling." *Strategic Management Journal*, 21: 1043-1059.

Rahmandad, H.

2005 "Three essays on modeling dynamic organizational processes." Massachusetts Institute of Technology.

Repenning, Nelson P.

2001 "Understanding fire fighting in new product development." *The Journal of Product Innovation Management*, 18: 285-300.

Repenning, Nelson P., and John D. Sterman

2002 "Capability Traps and Self-Confirming Attribution Errors in the Dynamics of Process Improvement." *Administrative Science Quarterly*, 47: 265-295.

Rivkin, J. W.

2000 "Imitation of complex strategies." *Management Science*, 46: 824-844.

2001 "Reproducing knowledge: Replication without imitation at moderate complexity." *Organization Science*, 12: 274-293.

Roberts, P. W., and G. R. Dowling

2002 "Corporate reputation and sustained superior financial performance." *Strategic Management Journal*, 23: 1077-1093.

Rosenbloom, R. S.

2000 "Leadership, capabilities, and technological change: The transformation of NCR in the electronic era." *Strategic Management Journal*, 21: 1083-1103.

Rudolph, Jenny W., and Nelson P. Repping

2002 "Disaster Dynamics: Understanding the Role of Quantity in Organizational Collapse." *Administrative Science Quarterly*, 47: 1-30.

Rumelt, R. P.

1991 "How Much Does Industry Matter." *Strategic Management Journal*, 12: 167-185.

Schreyogg, G., and M. Kliesch-Eberl

2007 "How dynamic can organizational capabilities be? Towards a dual-process model of capability dynamization." *Strategic Management Journal*, 28: 913-933.

Siggelkow, N.

2002 "Evolution toward fit." *Administrative Science Quarterly*, 47: 125-159.

Spanos, Y. E., and S. Lioukas

2001 "An examination into the causal logic of rent generation: Contrasting Porter's competitive strategy framework and the resource-based perspective." *Strategic Management Journal*, 22: 907-934.

Sterman, J. D.

1994 "Learning in and about complex systems." *System Dynamics Review*, 10: 91-330.

Sterman, J. D.

1989 "Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment." *Management Science*, 35: 321-339.

Sterman, J.K., et al.

1997 "Unanticipated side effects of successful quality programs: Exploring a paradox of organizational improvement." *Management Science*, 43: 503-521.

Sterman, John

2000 *Business Dynamics: systems thinking and modeling for a complex world*, 1 ed. Irwin: McGraw-Hill.

Taylor, T., and D. N. Ford

2006 "Tipping point failure and robustness in single development projects." *System Dynamics Review*, 22: 51-71.

Teece, D. J., et al.

1997 "Dynamic capabilities and strategic management." *Strategic Management Journal*, 18: 509-533.

Tripsas, M., and G. Gavetti

2000 "Capabilities, cognition, and inertia: evidence from digital imaging." *Strategic Management Journal*, 21: 1147-1161.

Wernerfelt, B.

1984 "A Resource-Based View of the Firm." *Strategic Management Journal*, 5: 171-180.

Williamson, O. E.

1999 "Strategy research: Governance and competence perspectives." *Strategic Management Journal*, 20: 1087-1108.

Winter, S. G.

2000 "The satisficing principle in capability learning." *Strategic Management Journal*, 21: 981-996.

2003 "Understanding dynamic capabilities." *Strategic Management Journal*, 24: 991-995.

Zollo, M., and S. G. Winter

2002 "Deliberate learning and the evolution of dynamic capabilities." *Organization Science*, 13: 339-351.

Tables and Figures

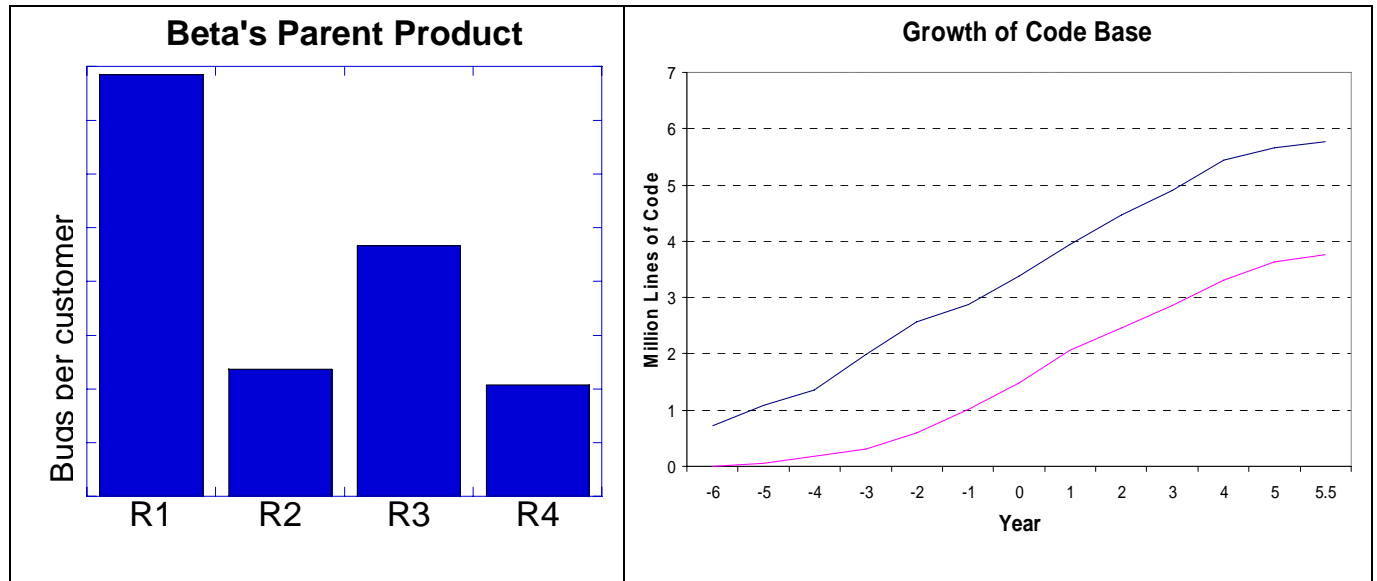


Figure 1- Beta's (parent product) quality (left) and size (right). Quality is measured in bugs reported per customer in the three months after the installation of a release. Scale is removed for confidentiality but is linear and starts from zero. Size is reported in new lines of code added to the code base, and year zero is when Alpha joined Sigma organization.

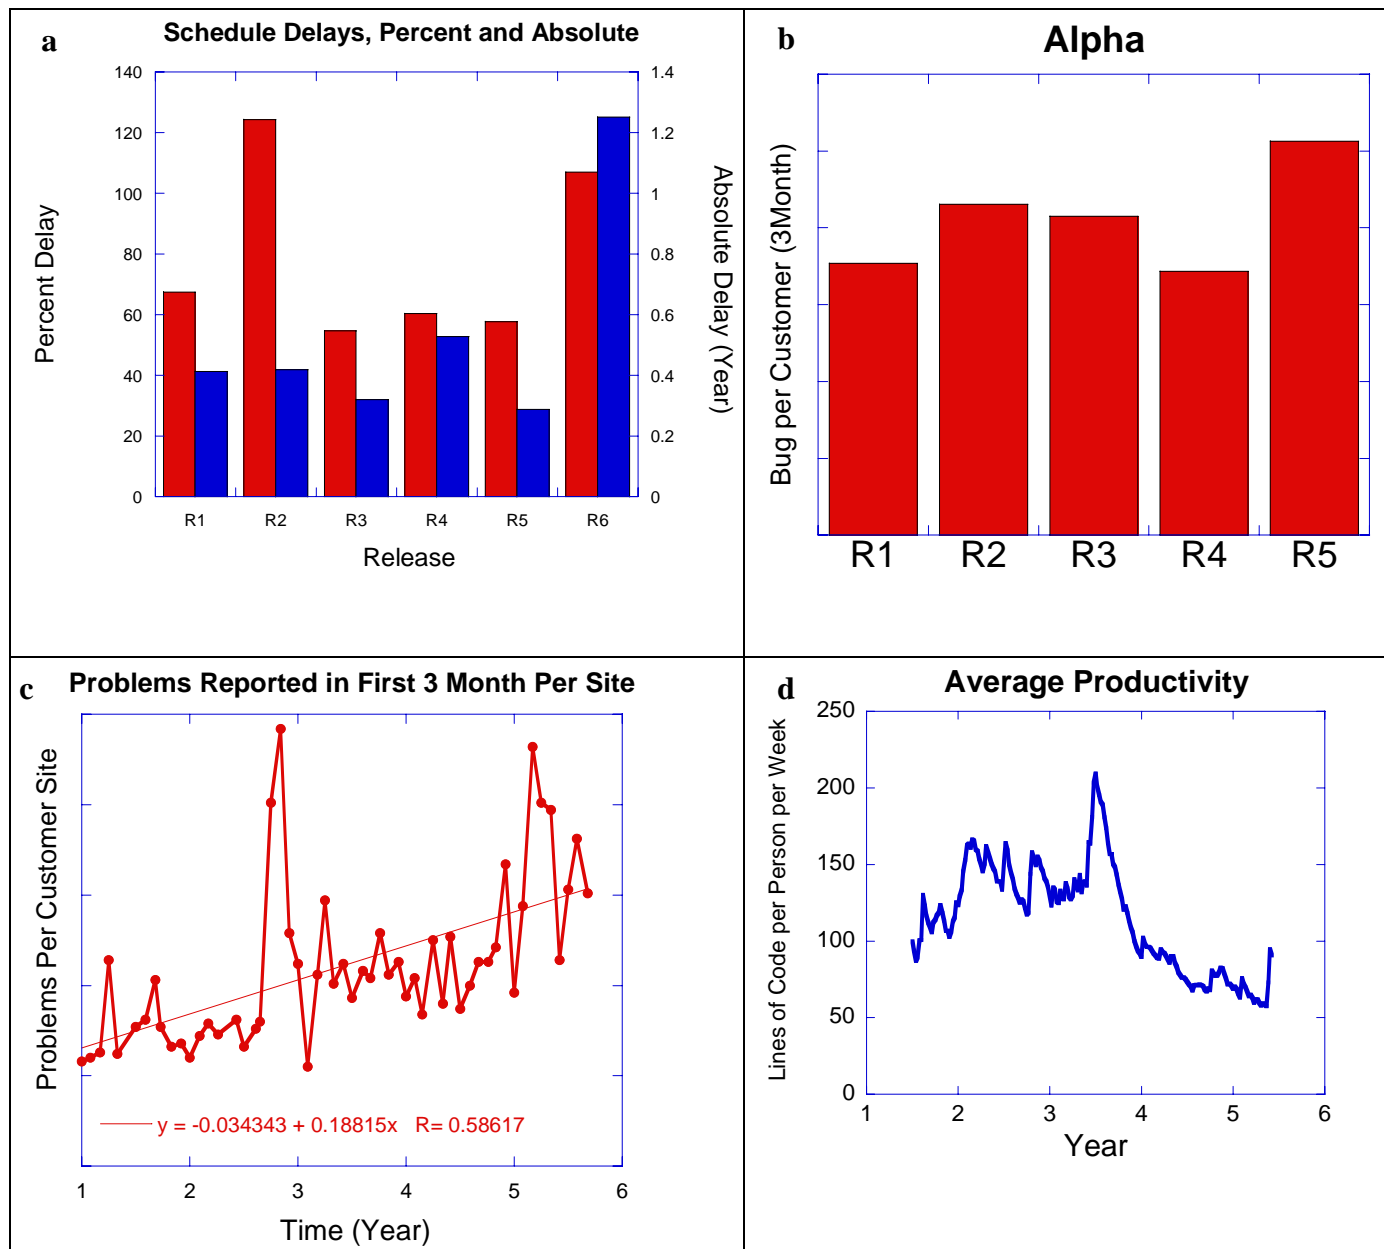


Figure 2- Alpha's performance on schedule (a), quality (b and c, both y axes start from zero and are linearly scaled, true scale hidden for confidentiality reasons), and productivity (d). a) The absolute (right bars) and percent (left bars) of delay for different releases of Alpha. Delays are expressed as the difference between the originally estimated lengths of the projects and their actual lengths and show no improvement pattern. b) The average number of bugs per customer in the first three month of installation, for different releases of Alpha. c) The same measure as b, but shown continuously. The slope of the curve is significantly positive, suggesting a gradual decline in quality. d) The productivity of Alpha team, expressed in the average number of lines of new code per person per week, shows a significant decline after year three.

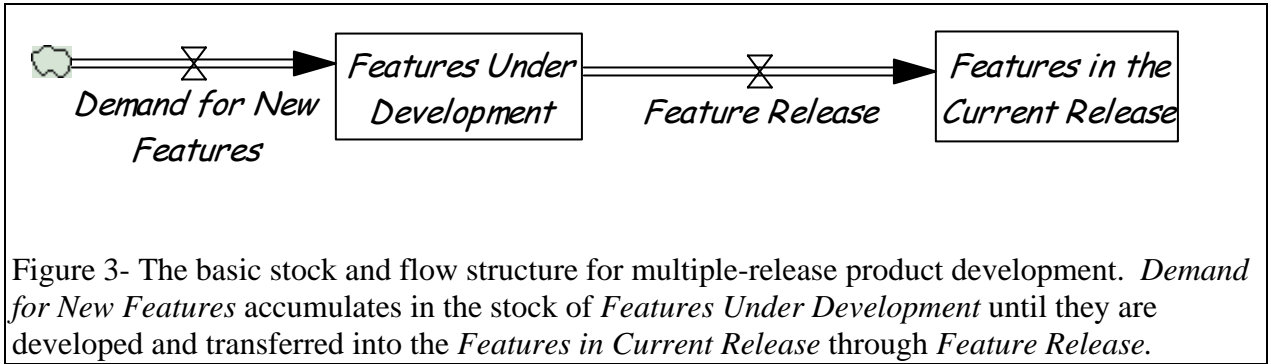


Figure 3- The basic stock and flow structure for multiple-release product development. *Demand for New Features* accumulates in the stock of *Features Under Development* until they are developed and transferred into the *Features in Current Release* through *Feature Release*.

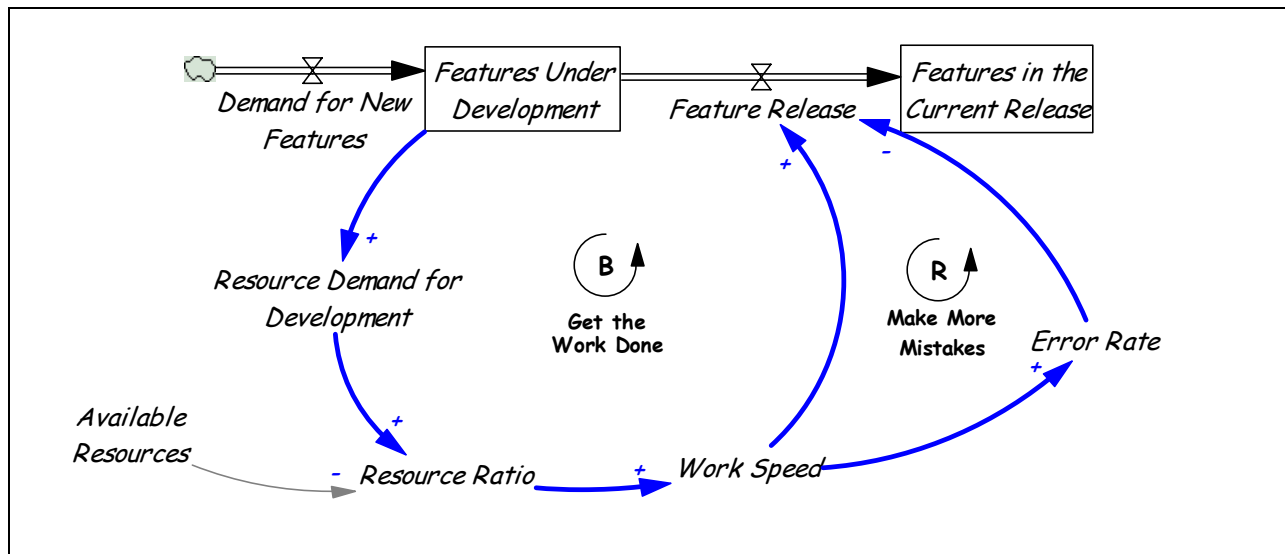


Figure 4- The basic structure of development process and the balancing (Get the Work Done) and reinforcing (Make More Mistakes) loops generated by considering the impact of resource ratio on work speed (gross productivity) and error rate.




Figure 5- a) Schematic inverse-U shape curve of Feature Release as a function of Resource Ratio. The picture highlights the stable equilibrium, the tipping point, and the maximum capacity. b) Increasing Feature Demand increases feature release (up to maximum capacity); it also reduces robustness, in terms of the shift in Resource Ratio needed to tip the system (the distance between stable equilibrium and tipping point). Full model formulations are in the online appendix.




Figure 6- The reaction of simulated PD system to a temporary demand shock under two slightly different shocks. a) The system recovers through the balancing process of “Get the Work Done” b) The system goes over the tipping point, the reinforcing loop “Make More Mistakes” dominates, and system performance deteriorates. Full model formulations are in the online appendix.

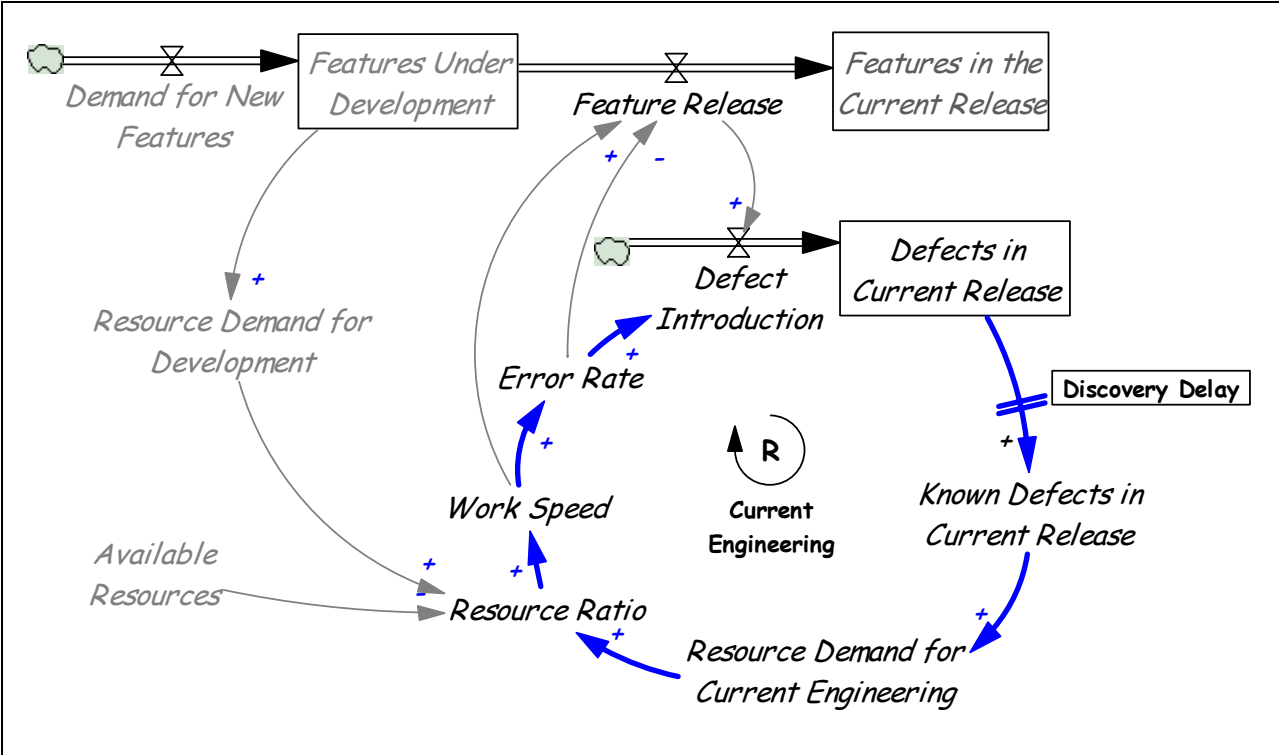


Figure 7- The structure of the model with the addition of products in the field. The new structures are highlighted through darker variable names and thick causal links.

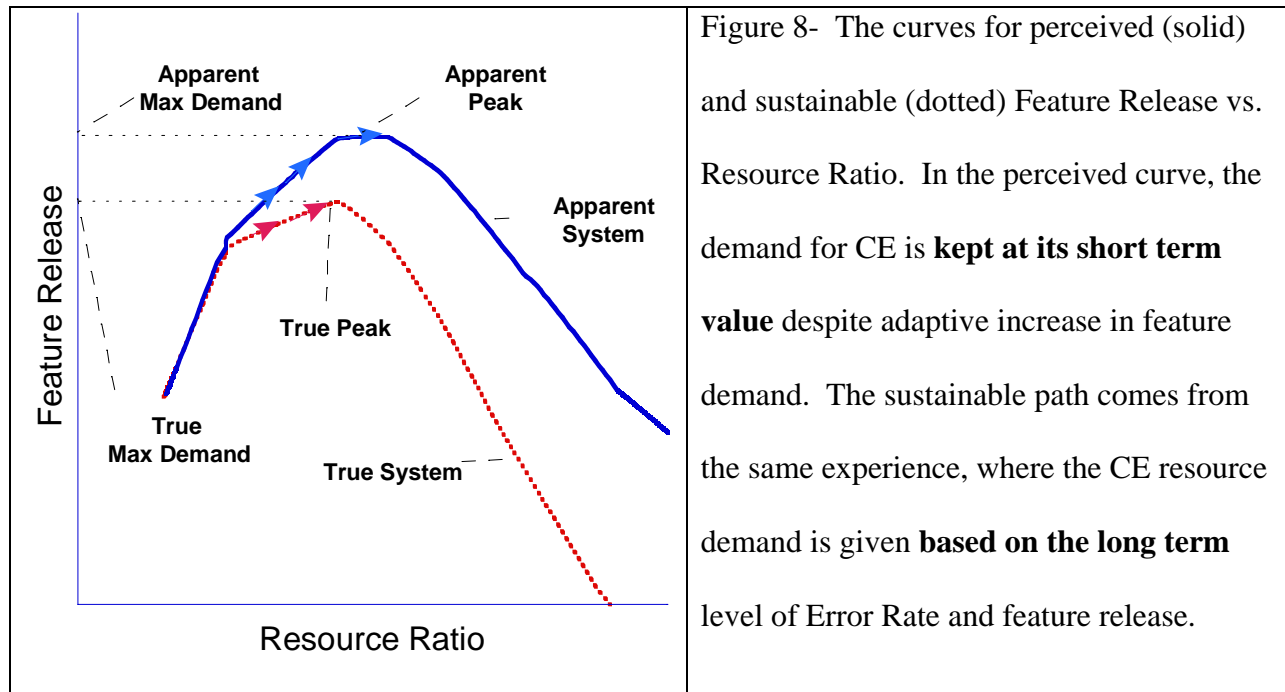


Figure 8- The curves for perceived (solid) and sustainable (dotted) Feature Release vs. Resource Ratio. In the perceived curve, the demand for CE is **kept at its short term value** despite adaptive increase in feature demand. The sustainable path comes from the same experience, where the CE resource demand is given **based on the long term** level of Error Rate and feature release.