

The Business School for the World®

Faculty & Research Working Paper

Explaining the Implementation Imperative: Why Effective Implementation may be Useful even with Bad Strategy

Eucman LEE Phanish PURANAM 2013/103/STR

Explaining the Implementation Imperative: Why Effective Implementation may be Useful even with Bad Strategy

Eucman Lee*

Phanish Puranam**

July, 2013

The authors thank Jerker Denrell, Thorbjorn Knudsen, Sangyoon Li, Oliver Bauman, Ulrik William Nash, Gael Le Mans, Hart Posen, Sendil Ethiraj and Chengwei Liu for helpful suggestions. Participants in the "First Cut" seminar series in Singapore provided very helpful feedback. Puranam acknowledges funding from the European Research Council under Grant # 241132 for "The Foundations of Organization Design" project.

* PhD candidate in Strategy and Entrepreneurship at London Business School, Strategy and Entrepreneurship Sussex Place, Regent's Park London NW1 4SA, United Kingdom. Email: <u>elee.phd2008@london.edu</u>

** The Roland Berger Chaired Professor of Strategy and Organisation Design, Professor of Strategy at INSEAD, 1 Ayer Rajah Avenue 138676 Singapore. Email: <u>Phanish.puranam@insead.edu</u>

A Working Paper is the author's intellectual property. It is intended as a means to promote research to interested readers. Its content should not be copied or hosted on any server without written permission from publications.fb@insead.edu

Find more INSEAD papers at http://www.insead.edu/facultyresearch/research/search_papers.cfm

Abstract

We propose a theory that explains why the relentless pursuit of effective implementation – the ability to get employees to take actions believed necessary to execute a strategy – may be useful even in a world in which the strategies being implemented are far from optimal. In keeping with typical models of organizational adaptation, we formulate a computational model in which an organization's strategy adapts based on performance feedback. However, the distinctive feature of our approach is that we abandon the "organization as a unitary actor" assumption, and model a separation of beliefs and actions. In our models, the strategy believed to be optimal by a manager is imperfectly implemented by an employee. The central insight is that when there is a separation of belief and action, effective implementation has benefits beyond the well-known effect of enabling exploitation of good strategies. It also enables the discovery of better strategies by allowing more effective learning from feedback on the value of current strategies.

1. Introduction

The links between strategy, organization and performance have always been central to the field of strategic management (e.g. Chandler, 1962). In this perspective, the path to superior performance lies in the implementation (or "execution" as it is sometimes called), through appropriate organizational arrangements, of a strategy that is consistent with the environment that the firm finds itself in (e.g. Lawrence and Lorsch, 1967; Galbraith, 1973; Nadler and Tushman, 1997; Hrebiniak, 2006). Thus effectiveness at strategy implementation- defined as the success with which strategic decisions are translated into appropriate organizational arrangements (Neilson, Martin and Powers, 2008)- is generally seen as a key ingredient for successful organizations (Hrebiniak, 2006; Johnson 2004; Neilson, Martin and Powers, 2008; Noble and Mokwa, 1999; Yang, Guohui and Eppler, 2008). However, it is not easy to achieve, as lower level managers may not execute strategies perfectly because some degree of employee discretion in actions is inevitable (Barnard 1938; Simon 1951), and no organization design can completely control behavior (Galbraith and Kazanjian, 1988; Govindarajan, 1988). Yet, the acknowledged difficulty of achieving it does not appear to have diminished the desirability of effective strategy implementation in the world of managerial practice (see Noble and Mokwa, 1999; Hrebiniak, 2006; Yang, Guohui and Eppler, 2008; Bossidy and Charan, 2002).

The dominance of this "implementation imperative"- *the view that perfect implementation of strategy is desirable and must be relentlessly pursued* - in managerial discourse (see for instance Hrebiniak, 2006; Neilson, Martin and Powers, 2008; Bossidy and Charan, 2002) is puzzling on at least two counts: It appears to repose too much faith in the quality of the strategies imposed "top down" on the organization, and too little in the value of "bottom-up" exploration for better strategies. It is well recognized today, even by its proponents, that top-down strategic planning is

hardly infallible and must constantly confront the challenge of inadequate local knowledge (e.g. Breene et al 2007; Kaplan and Norton, 2005). Indeed, a key insight from the perspective on organizations as adaptively rational systems (e.g. Simon, 1947; March and Simon, 1958; Cyert and March, 1963) has been that imperfect implementation of strategy may not be such a bad thing in a world where strategies themselves may often be incorrect; imperfect implementation may in fact be a valuable source of "bottom-up" exploration for better strategies (Burgelman, 1983; March, 1991; Levinthal and March, 1993; Eisenhardt and Brown, 1999; Benner and Tushman, 2002; 2003).

In the process of the search for better alternatives, variance in actions generates explorationthe trying out of alternatives with unknown value rather than alternatives with known valueswhich serves a useful purpose by making it more likely that better alternatives can be found. Excessive exploration however prevents exploitation of the valuable alternatives already known, giving rise to the famous exploration-exploitation tradeoff in learning systems (Gittins,1979; March 1991; Sutton and Barto, 1998). Thus, some degree of variance in realized organizational actions around the intended action may be a valuable source of exploration for better strategies (March, 1991), rather than something to be suppressed entirely.

To make these ideas concrete, consider the phenomenon of "bootleg innovation" in which individuals ignore top-down management directives, to make their own choices on how to invest company resources to pursue innovation (Knight, 1967). These efforts are in direct contravention of official strategy and so represent a failure of strategy implementation. Yet, it is now welldocumented that such deviations from strategy are often beneficial to the firm (Augsdorfer, 2005). Iconic cases include the story of Apple's Graphic calculator which was developed by two employees who worked on it without official authorization, and later despite being asked to desist; in fact they are known to have sneaked in to the premises to continue working on it even after they were no longer employed by Apple. Eventually, the product was recognized as a valuable add on to the Mac OS.² Similarly, membrane filtration techniques for the production of pharmaceutical compounds are in wide use today, but were initially the result of bootlegging activity by scientists at Beecham (now Glaxo Smith Kline) in the 1980's, who continued working on it despite official orders to stop (Augsdorfer, 2005). Delegation in resource allocation decisions to functional and middle managers may produce a similarly useful divergence between intended and realized strategy in complex multi-layered organizations (Bower, 1982; Burgelman, 1983).

Even when these bottom-up deviations from official top down strategy are not as explicitly intended as in bootlegging activities or delegated resource allocation, natural variations and even errors can often lead to improvements in organizational knowledge. The history of science and innovation is replete with instances of serendipity, which by definition must have arisen through deviations from plans (e.g. Merton and Barber, 2003; Nonaka and Takeuchi, 1995). Too rigid an emphasis on implementation of existing imperfect strategies may well curtail such useful variation. For instance, Benner and Tushman (2002) have documented the adverse consequence for exploratory innovation of process management techniques (such as total quality management) that are explicitly aimed to reduce variation and improve alignment between strategy and employee actions (also see Benner and Tushman, 2003; Slater, Hult and Olson, 2010). Edmondson (2008) has argued explicitly that too much focus on execution- as-efficiency crowds out opportunities for execution-as-learning.

² http://mentalfloss.com/article/31077/when-his-project-was-canceled-unemployed-programmer-kept-sneaking-apple-finish-job

The insight that some degree of imperfect implementation could well be desirable in a world in which the optimal action is rarely known ex ante but must be discovered through trial and error, also has very robust analytical foundations in the formal modeling literature on learning (e.g. Gittins, 1979; Sutton and Barto, 1998; March, 1991). It has also become quite deeply embedded in the management literature on organizational learning and innovation, with an extensive literature developing around the exploration-exploitation tradeoff as well as proposed organizational solutions to manage it better (e.g. Levinthal and March, 1993; Tushman and O'Reilly, 1997;Benner and Tushman, 2003; Gibson and Birkinhaw, 2006; Gulati and Puranam, 2009). Yet, it appears to have made few inroads into the world of strategy implementation. Senior managers continue to list improving strategy execution as their top priority³, and the practitioner literature offers many exhortations to pursue the ideal of perfect implementation of strategy (Sull and Spinosa, 2007; Neilson, Martin and Powers, 2008; Bossidy and Charan, 2002).

In this paper, we construct a formal computational model that helps us to understand why effective implementation may be desirable *even* in a world in which there exist potential gains from bottom-up exploration and there are limits to the quality of top-down strategies. The resulting insights suggest an explanation for why the implementation imperative continues to feature so prominently in the discourse of management practice – it offers an argument for the "rationality of practice", as it were. In keeping with typical models of organizational adaptation (Cyert and March 1963; Herriott, Levinthal et al. 1985; March 1991; Denrell and March 2001; Posen and Levinthal, 2012), we formulate a model in which an organization's strategy adapts based on performance feedback. Thus, the organization in our model is assumed to be adaptively rational- initial beliefs about the right strategy may be inaccurate to varying degrees, but adapt on the basis of feedback. In this sense, the model is a classic model of adaptive rationality through

³<u>http://www.conference-board.org/publications/publicationdetail.cfm?publicationid=1746</u>

experiential learning, as it features belief and action revision on the basis of feedback on current actions (Thorndike 1998; Sutton and Barto 1998).

However, the key distinction is that we enrich current models by explicitly incorporating the feature that the agent who holds beliefs and the agent who executes them are not the same. This recognizes the property that organizations are multi-agent hierarchical systems in which delegation is pervasive (Simon, 1947; March and Simon, 1958). In the model, a manager holds a belief about the appropriate strategy for the organization at a point in time, and a subordinate implements this belief. The manager then updates beliefs about the appropriate strategy based on feedback about organizational performance. Our model helps us to explore the general intuition that the process of learning from feedback should be more complicated when beliefs and actions are separated across agents in a multi-agent adaptive system, in a precise and rigorous way. Specifically, we study the effects of imperfect communication between the manager and the subordinate, imperfect observation of the sub-ordinates actions by the manager, and top down exploration by the manager on the value of bottom up exploration by the subordinate. Our results provide insights on the conditions under which perfect implementation by the subordinate is desirable, even if the strategies the manager wants the subordinate to implement are far from perfect.

2. Model

Since the canonical representation of the exploration–exploitation tradeoff across multiple literature streams is based on learning in the multi-arm bandit task (Gittins, 1979; Sutton and Barto, 1998; Denrell and March 2001) we explore a model that uses the same task but features a separation of beliefs and actions across multiple agents. Posen and Levinthal (2012) provide a recent detailed introduction to the applications of the bandit model in organization theory and related literatures. We therefore concisely describe the model in terms of its four basic components: the task environment, the agent's representation of the task environment, a process for transforming the agent's representations (i.e. learning), and a choice process through which the agent selects actions within the representation.

The *task environment* of the *n* armed bandit model takes its name from analogy to slot machines in casinos. In the casino, a gambler faces a slot machine with *n* arms each with unknown underlying payoff distributions. In organizational applications, the arms may represent different strategies, projects, or more generally choices with unknown expected payoffs and variances. As Posen and Levinthal (2012) note, the primary feature of this problem is that information about the value of an arm can only be obtained by selecting it. Thus at any point in time, an organization confronting such an *n* armed bandit must choose one of the arms, which results in a payoff that is a realization from a (possibly arm specific) probability distribution. These payoffs constitute visible performance measures for organizations – such as profits, market share or market capitalization.

We assume that our organization consists of a manager (M) and her subordinate, (S). In our model, M formulates strategies, communicates them to S (possibly imperfectly) and observes how S implements them (possibly imperfectly). S implements the strategy as understood by him, but imperfectly; this can be understood recursively to be the case because of S's own subordinates who implement S's intentions albeit with errors. Thus a natural interpretation of M and S might be that of a CEO and Division general manager in a multi-divisional corporation.

Both agents, at any point in time may have a *representation* of the task environment in the form of beliefs (not necessarily knowledge) of the task environment, in terms of an associated set of estimated payoffs for each arm. Following prior work, we assume that for each arm $a \in$

 $\{1,2,...,n\}$, there exists an unknown true expected reward $Q^*(a)$ and based on the accumulated reward information, agent j's estimate of this (where j={M, S}), $Q_{it}(a)$ at time *t*, is calculated as below

$$Q_{jt}(a) = \frac{1}{k_a} \cdot \sum_{i=1}^{k_a} r_i \tag{1}$$

The vector of $Q_{jt}(a)$'s thus constitutes the agent j's representation of the task environment. Here, r_i represents the realized reward of the action a at i^h trial and $Q_{jt}(a)$ is the sample average after it has been <u>believed</u> to have been tried k_a times. (The fact that k_a is a belief will become critical when we introduce separation of belief and action into the model). This rule for constructing the estimate comes from the assumption that the estimate at any point in time of an arm's value is the average of the entire history of rewards arising from having selected that arm (March, 1996; Posen and Levinthal, 2012). This is how the agent's representation of the task is *transformed* through feedback. Note that (1) implicitly states that the estimate of the value of a strategy increases if its current performance exceeds the aspiration level defined by historic average, else it decreases.

To complete the specification of the model, we need to specify a *choice* process. We assume a standard choice process from the literature on bandit models. The probability of an agent j choosing action a at time t is calculated using the softmax rule as below, where n is the total number of alternative actions that the agent confronts.

$$P_{jt}(a) = \frac{exp\left(\frac{Q_{jt}(a)}{\tau_j}\right)}{\sum_{i=1}^{n} exp\left(\frac{Q_{jt}(i)}{\tau_j}\right)}$$
(2)

 $\langle \mathbf{a} \rangle$

This softmax action selection rule has been found to effectively describe how humans select actions in trial and error learning situations (Camerer and Ho, 1999; Daw et al, 2006), and has also been used to model organizational adaptation processes (e.g. Posen and Levinthal, 2012). High values of agent specific $\tau_j \gg 0$ will promote exploration in the sense that the agent becomes more likely to select actions other than those currently believed to be the best, while low τ_j results in the current best belief to be selected.

The dynamic model represented by equations 1 and 2 capture a reinforcement learning process - whose variants are known under the labels of trial and error learning, experiential learning, operant or instrumental conditioning or "win-stay-lose-shift" rules in the relevant literatures in psychology, computer science, organization theory and evolutionary biology (Thorndike, 1911; March, 1991; Domjan 2010; Nowak and Sigmund, 1993; Sutton and Barto, 1998).

The sequence of actions by the two agents M and S in the model can now be specified:

<u>Step 1</u>

M inspects her representation and selects a strategy (arm) based on (2). In our model, each arm represents a strategy- possibly a detailed set of instructions on what actions S should take to implement that strategy. M's beliefs may be erroneous in the sense that they may correspond only approximately to the real strategy which is necessary to maximize organizational performance given the particular environmental conditions. This is either because the beliefs are coarse, inaccurate, or both. The gap between underlying reality and the beliefs of the manager reflects imperfect knowledge of the manager. M may conduct some degree of *top-down exploration*; replacing τ_j with τ_M in (2), as τ_M increases M is more likely to choose a strategy other than what appears the best in M's current representation. M then communicates this choice

 a_{Mt} to S. In a very abstract form, we assume that the manager's orders are conveyed by setting incentives and exerting authority to motivate and direct the desired actions of the subordinate.

<u>Step 2</u>

S receives the choice of M, and updates her own representation by linearly combining its own representation and M's order: $Q_{St}(a_{Mt}) = \lambda \cdot Q_{St} + (1-\lambda) \cdot O(a_{Mt})$. The parameter λ represents the degree of *imperfect communication* of strategy. When $\lambda=0$, M has been able to communicate her desired strategy to S effectively, and S's new representation (Q_{St}) reflects this by increasing the estimated value for the action that M ordered (a_{Mt}) by an Order function (O) that increases the attraction of the corresponding arm by unity (the maximum expected payoff for any arm). For instance in the case of a two-armed bandit, let's take the case that $Q_{St}=[0.5, 0.5]$ while $O(a_{Mt})=[1,0]$, with $\lambda=0.5$. Then subordinate's updated attraction $Q_{St}=0.5\times[0.5, 0.5] + 0.5\times [1,0]$ = [0.75, 0.25]). Conversely when $\lambda = 1$, S effectively has not "heard" M at all. An alternate interpretation of λ , is in terms of the *delegation* of strategy making by M to S. When it is high, M lets S independently choose the strategy to implement based on S's own knowledge. When it is low, M instructs S to implement M's preferred strategy.

<u>Step 3</u>

S attempts to implement the strategy as understood by him, incorporating M's orders (if $\lambda < 1$) using (2). Note that S's realized actions a_{St} is a function of her perception of M's dictated strategy, but also features some divergence. These deviations could be thought to arise from lack of sufficient specificity of the understood strategy, or through insufficient motivation for the employee to take the desired action, or because of imperfect implementation of S's choices by S's own subordinates. This divergence is inversely related to the implementation competence of the organization. Thus, in an organization with low implementation competence the

subordinate S may enact an unintended action with some positive probability even if the strategy as understood by him does not require him to do so.

Replacing τ_j with τ_S in (2), we use the parameter $1/\tau_S$ to operationalize the *implementation competence* of this organization: When τ_S is low, the subordinates actions conforms to the strategy as she understands it; as τ_S increases, the quality of implementation falls, and the patterns of sub-ordinate choices essentially become random and no longer reflect the strategy as he understands it.

<u>Step 4</u>

The payoff of S's selected arm a_{St} is realized. M and S update their representations using (1). While S knows its own selected actions perfectly, M observes the selected arm with probability of *observation error* μ , such that with this probability M thinks the selected action was one of the n arms at random. The parameter μ may be thought of as capturing the accuracy of measurement of implementation.

The resulting model now allows us to examine how this simple 2-agent organization performs in a task environment that resembles a multi-armed bandit task (i.e. with discrete strategic alternatives whose value can only be learnt through trying them), in which both bottom-up exploration (τ_5) as well as top-down exploration (τ_M) are feasible, and in which top managers may be unable to specify or observe the actions of subordinates (via λ and μ). Our experiments with the model involve varying these four parameters in systematic ways to observe the impact on organizational performance. While we primarily conceptualize and discuss the model in terms of strategy implementation, the analytical structure more generally captures any process of experiential learning with separation of beliefs and actions.

3. Analysis & Results

We operationalize this dynamic model by setting up a n=10 armed bandit task. The true underlying reward associated with each arm is drawn from a beta distribution with both $\alpha = \beta = 2$. The beta distribution has been used in prior simulation studies as it has similar properties to the Gaussian distributions, but we can control its boundaries. The realized reward when an arm is actually selected in a trial is the true expected reward of that arm plus a noise term drawn from a normal distribution with zero mean and unit variance. The initial estimate of each arm's value held by M and S is assumed to be uniformly distributed as zero. The Performance for an organization in any period is simply the cumulative received payoff up to that period. To illustrate the underlying processes effectively, we also plot the exploration propensity and knowledge level (Posen and Levinthal, 2009) of the organization. The Exploration propensity is the proportion of past periods in which the organization changed its strategy relative to the prior period. Knowledge level is the probability that in a given period, the strategy believed to be the best (i.e. has the highest expected payoff) is indeed the true best strategy. We report results for performance, exploration propensity and knowledge level averaged across 1000 organizations over 500 periods.

Scenario 1: Baseline case with no separation of beliefs and actions

As a baseline, we first reproduce the standard result on the exploration –exploitation tradeoff with no separation of belief and action. We assume that $\tau_M \sim 0^4$, $\lambda=0$, and $\mu=0$. While the organization does consist of two agents, M and S, these settings reflect a case where the two effectively act as one- there is no error in communication or observation between M and S, and the only source of exploration in the organization is the result of imperfect implementation by S

 $^{^{4}\}tau$ cannot be zero in softmax strategy. Therefore, we used $\tau = 0.01$ instead of zero as it is the minimum value we can use for computations.

of M's orders. This is equivalent to the organization being a unitary actor, with S serving as a perfect pair of arms and eyes for M; there is effectively no separation of beliefs and actions because M's beliefs are transmitted without error to S, and S's actions are transmitted without error to M.

Figure 2a shows how organizational performance changes over increasing levels of bottomup exploration (τ_s). The vertical axis represents cumulative performance over five-hundred periods (without time discount) and the horizontal axis represents the inverse of implementation competence (τ_s). For the moment, we will focus on the line showing organizational performance for $\mu = 0$ (Figure 2a reports results for this and the next set of experiments together to conserve space). This inverted-U shaped line with an optimal $\tau_s = 0.5$ replicates the well-known exploration and exploitation trade-off: a moderate level of subordinates deviation from strategy delivers superior long term performance as it brings in a chance of probing unknown alternatives at the cost of not doing what they ought to be doing. Figure 2b helps to understand the underlying tradeoff more clearly. Decreasing implementation competence or equivalently, increasing bottom up exploration leads naturally to greater propensity to explore, as well as improvements in knowledge from exploration. However, beyond a certain level of exploration, improvements in knowledge cannot be utilized, because implementation of the strategies based on this knowledge is difficult. This generates the internal optimum for τ_{S} . While we plot the knowledge level for M, the knowledge level for S is the same in this setting, because M perfectly observes S's actions and updates her own beliefs in the exact same way as S does.

These results establish the baseline intuition that the pursuit of perfect implementation is unlikely to be optimal in a world that has the features of the multi-arm bandit task we have set up- where the best strategy is unknown ex ante and the value of strategies must be discovered through trial and error. This underlines the point that when the organization effectively has no separation of belief and action (because λ =0, and µ=0), and imperfect implementation is the only source of exploration by the organization ($\tau_M \sim 0$), then perfect implementation is not optimal, because bottom up exploration is valuable (i.e. $0 < \tau_5$). As March (1991:7) is widely cited for noting, organizations that "engage in exploration to the exclusion of exploitation are likely to find that they suffer the costs of experimentation without gaining many of its benefits," whereas organizations that "engage in exploitation to the exclusion of exploration are likely to find themselves trapped in suboptimal stable equilibrium." For an organization with no separation of beliefs and actions, some degree of imperfect implementation is therefore useful in a world in which initial knowledge is poor.

Scenario 2: Observation errors in the measurement of implementation success

In scenario 1, we modeled the organization as acting effectively like a unitary actor, as there was no separation between beliefs and actions, even though these were held by two different agents. We now introduce the first element of separation between beliefs and actions, by considering a scenario in which the manager sets strategy and communicates this perfectly to the subordinate, who executes the strategy imperfectly (as before); however now the manager observes the subordinates actions but with a probability of erroneous observation (μ >0).

This scenario focuses on the challenge of being able to measure the success of implementation (i.e. whether subordinates did what they were told to) independent of organizational performance. There are a number of reasons why in hierarchical organizations, superiors may be unable to perfectly observe the actions of subordinates. These include distinct bases of knowledge, distinct physical locations, different timescales of operation as well as the sheer cognitive burden on a superior in a hierarchical structure faced with the prospect of

observing a large number of direct and indirect subordinates. Yet, imperfect means to observe a subordinate's actions may still be possible. Balanced scorecards, key performance indicators and management by objectives allow measurement of implementation success, albeit imprecisely. Interestingly, performance measures of lower level organizational units can correspond to implementation performance as seen from higher levels of aggregation. For instance, internal targets for production cost or customer service lead time are measures of the successful implementation of strategies aimed at lowering costs or improving customer willingness to pay. The variability in the existence and precision of these measures is captured by μ .

In this set of experiments, we keep other variables $\tau_M \sim 0$ and $\lambda=0$. Thus we are investigating a single consequence of separating belief from action (between M and S respectively) – that M can no longer know with precision what actions S actually undertook. Obviously, if there was no separation of belief and action, and the organization was a unitary actor, then this problem would not arise. At the same time, we continue to assume that there is no communication problem between M and S, and that M does not introduce any top-down exploration.

Figure 2a also demonstrates the effects of observation errors (μ) on cumulative organizational performances over various levels of bottom up exploration (τ_S). There are two key aspects of these results worth noting. First, as observation error (μ) increases, the optimal level of implementation error (τ_S) moves toward zero. For sufficiently large observation errors, there is a strong implementation imperative (i.e. it is better to improve implementation). Second, when implementation competence is high (i.e. there is no bottom up exploration, $\tau_S \sim 0$), then optimally observation error is greater than zero.

To understand the intuition behind these results, it is useful to turn again to a graph of the propensity to explore and the knowledge level (Figure 2c). The first thing to note is that

observation error itself is a source of exploration for the organization. By decoupling the link between received payoff and the probability of selecting the same action again, observation error induces a degree of exploration. It is modest because it decreases the probability of selecting the same action after a good outcome (an increase in exploration), but may also increase the probability of taking the same action after a bad outcome (a decrease in exploration). To see this consider a case where S selects action 1, and a good payoff is received (i.e. a payoff above the current value estimate of that action). In the absence of observation error, this would raise the value estimate of that action and make it more likely to be selected again. In the presence of observation error, because the received payoff is likely to be attributed to some other action, this is less likely; therefore the organization de facto explores. Conversely, if the payoff was bad (i.e. below current value estimate of the action), in the absence of observation error this action would be less likely to be selected in the next round. However, with observation error, because the payoff is attributed to some other action, there is a higher (but still small, as action 1 is just one of m-1 actions that has "escaped" being attributed the poor performance) chance that the same action is selected again (which indicates a decrease in exploration).

The second consequence of observation error is on the knowledge of M. Observation errors interfere with the learning process represented by equation (1) because the values of both a and k_a held by M will be erroneous; in effect, errors in updating of beliefs is introduced (through μ), which will result in a bias in the estimated value of the action. This is because with observation error, by definition, there is a non-zero probability that the estimated value of a selected action is incorrect; even if the action to which the realized payoff is erroneously attributed is random, the lack of correct update to the selected action is systematic. As a consequence, as observation errors increase, there is a decrease in knowledge level for M (Figure 2c). To contrast this with

the consequences of bottom up exploration, compare with Figure 2b and with Figure 2a when τ_m =0; excessive bottom up exploration lowers performance but not knowledge levels. Put differently, implementation errors create exploration but not bias in beliefs; observation errors create exploration and bias. Figure 2d shows optimal (highest performing) combinations of implementation errors and observation errors; they are substitutes in organizational performance, but the marginal rate of technical substitution is not one.

The implications of these results are that in organizations with a high degree of implementation competence, a modest level of observation errors could usefully contribute a degree of exploration. As can be seen in Figure 2a, when there is no bottom up exploration, organizational performance is highest when observation error is μ =0.3. Indeed upto a value of τ_s <0.3, adding observation error always increases performance. However, at lower levels of implementation competence (i.e. higher levels of τ_s) adding observation error diminishes performance. One reason for this is simply too much exploration – if the organization was already close to the optimal degree of exploration, then increasing exploration must harm performance. The other is that as observation error increases, knowledge levels diminish. Thus as observation errors increase, the implementation imperative strengthens, as optimal bottom-up exploration falls to zero.

Therefore we can conclude that the implementation imperative is likely to be strong despite potential gains from bottom-up exploration when it is difficult for senior managers to observe exactly what the subordinates have done- in other words, when it is difficult for them to assess the effectiveness of implementation. This is the first way in which a separation of beliefs and actions can lead to an implementation imperative, despite all the usual arguments for the benefits of exploration in a world of poor initial knowledge (March, 1991; Sutton and Barto, 1998).

Scenario 3: The effect of imperfect communication of strategy

So far we have assumed that the strategy of the top manager is communicated perfectly to the subordinate- the challenge lies in implementing it. While it may be tempting to view imperfect communication as part of the problem of implementing strategy, we believe there is a gain in analytical clarity by separating the two; after all it is possible that a poorly communicated strategy is perfectly implemented.

In the next set of experiments, we investigate the effects of imperfect communication of strategy by M to S (i.e. λ). The other variables are kept at $\tau_M \sim 0$ and $\mu=0$. This corresponds to an organization in which we have introduced a second way in which a separation between beliefs and actions may exist, by considering a scenario in which the manager sets strategy, the subordinate executes the strategy imperfectly (as before); however now while the manager observes the subordinates actions perfectly ($\mu=0$), he may be unable to communicate her selected strategy to the subordinate ($\lambda>0$).

It is a recognized challenge for senior level managers to effectively communicate their intended strategy to lower level managers charged with implementing the strategy (Alexander, 1985). In part this challenge arises from the somewhat abstract nature of strategy itself, but more critical, it is the consequence of vertical specialization. Managers at different layers in a hierarchy know and are capable at different things. While lower level managers may understand their world in terms of fairly detailed if narrow scope representations, senior managers necessarily utilize more abstract and broader representations. These distinctive "thought worlds" they inhabit also imply communication difficulties between them (Dougherty, 1992). While senior managers may be able to convey the general sense of their intentions, there may be

significant gaps between how lower level managers understand these intentions in the operational terms they are accustomed to dealing with. The importance (as well as the difficulties) of being able to communicate the organization's strategy therefore take center stage in managerial discourse on strategy implementation (Bossidy and Charan, 2002). The parameter λ in our model captures the extent to which this is achieved; low values indicate relatively unambiguous communication of strategy in a manner that the subordinates understanding of the managers preferred strategy is high; high values indicate that the subordinate's understanding of the managers preferred strategy actually bears no connection to the manager's preferred strategy.

The results of this set of experiments are reported in Figure 3a. There are again two key aspects of these results. First, as the extent of imperfect communication (λ) increases the optimal level of implementation error (τ_S) moves toward zero. In other words, as the difficulty of communicating stagey from M to S increases, the stronger the implementation imperative (i.e. there are incentives to improve implementation). Second, when implementation competence is high (i.e. there is no bottom up exploration, $\tau_S \sim 0$), then communication error does not appear have significant impact on organizational performance.

To understand these results, as before it is useful to turn again to a graph of the propensity to explore and the knowledge level (Figure 3b), for the case when there is no bottom-up exploration as well as the optimal bottom-up exploration when there is no constraint on communicating strategy. It is clear that imperfect communication of strategy, by itself, is not a source of exploration for the organization. However, when there is bottom-up exploration, it is leveraged by imperfect communication.

The intuition can be understood as follows: Consider first the case when S engages in bottomup exploration. When M's selected strategy is conveyed perfectly to S, it effectively makes that strategy much more salient (i.e. with a higher value estimate) than even in M's own representation. For the same action value estimates, S would have to explore more when $\lambda=0$ to escape making a greedy choice (the one ordered by M). Perfect communication of strategy effectively dampens S's exploration. As the degree of imperfect communication increases, it lowers the weight placed on M's preferred choice in S's representation. Therefore it allows for greater exploration by the organization, for the same level of bottom-up exploration by S.

Now consider the case, when neither M nor S engage in any exploration, then the organization as a whole engages in greedy search; the first choice with a payoff that is better than the estimates of the rest (recall we start all estimates at zero), will be the choice picked in perpetuity. More importantly, *both M and S will always prefer the same choice*, which S will pick in a greedy manner. Increasing the degree of imperfect communication of strategy does not affect this result, because all that imperfect communication does, when neither M nor S explore is to change the weight that S places on M's preferred choice in her own representation; but because neither explore, their representations are identical anyway. For this reason, imperfect communication leverages bottom-up exploration by S, but does not contribute any exploration by itself- it undoes the dampening of exploration that would have occurred with perfect communication. This is a significant difference between the consequences of observation error (Scenario 2) and imperfect communication.

A second important distinction between the consequences of observation error and imperfect communication can also be seen in Figure 3b; the effect of imperfect communication on the knowledge of M is benign- it does not harm knowledge level in any way. In contrast, as we saw in Scenario 2, observation errors introduce bias in the beliefs of M. Note that these consequences are not necessarily because of the hierarchical relationship between M and S; rather they are because of a separation of belief and action between M and S, which could arise even if they were peers. Figure 3c shows optimal (highest performing) combinations of imperfect communication levels and bottom-up exploration; as communication of strategy is clearer, the optimal level of bottom-up exploration must increase for the organization as a whole to explore sufficiently.

The implications of these results are that the organization in fact benefits from less than perfect implementation when strategy is communicated effectively by senior managers to their subordinates; put differently bottom-up exploration is valuable. Conversely the implementation imperative is likely to be stronger when it is *harder* to communicate strategy effectively from top managers to their sub-ordinates. Further, our results suggest that the inability of top managers to clearly communicate strategy may be less of a problem both for organizational performance as well as their own understanding of the optimal strategy, than the inability to observe how strategy was implemented in the sense that both organizational performance and the level of knowledge is relatively less sensitive to the top manager's ability to communication strategy (See Figures 3c vs. 2d, and 3b vs. 2c).

Therefore we can conclude that the implementation imperative is likely to be strong despite potential gains from bottom-up exploration when it is difficult for senior managers to communicate exactly the strategies they want implemented. This is the second way in which a separation of beliefs and actions can lead to an implementation imperative, despite the usual benefits of exploration with poor initial knowledge.

Scenario 4: The effects of simultaneous top-down and bottom-up exploration

So far we have assumed that M does not explore- only S engage explicitly in explorative actions (as a consequence of imperfect implementation). In the next set of experiments, we examine the consequences of top-down exploration by M on the optimal value of bottom-up exploration by S. Top-down exploration could correspond to investment in strategic analysis, hiring consultants, strategic planning processes etc. These represent ways in which M is motivated to engage in non-greedy selection of strategy, to be conveyed to S for implementation.

We first hold λ =0 and μ =0 so that we can focus purely on the interaction between bottom-up and top-down exploration. This corresponds to an organization in which both top and bottom up exploration are taking place, but there are no observation errors or imperfect communication of strategy. **Figure 4a** shows that when λ and μ are zero, the optimal degree of bottom-up exploration declines with top down exploration. This is a fairly intuitive result, as in effect the organization over-explores when there is both top-down and bottom up exploration, so that if M is engaging in exploration, then there is an implementation imperative. **Figure 4b** shows that the level of exploration increases significantly as top-down exploration occurs in addition to bottomup exploration, though there is no perceptible improvement in knowledge level. **Figure 4c** shows optimal combinations of top-down and bottom-up exploration; it is clear that they are substitutes, though a higher level of exploration by the subordinate than the manager is necessary for a given level of performance to be achieved. This is because when M does not explore and there is no mis-communication, as we have seen in the results of scenario 3, the sub-ordinates exploration is effectively dampened down.

In sum, these results show that bottom-up exploration through imperfect implementation can be a substitute for top down exploration, that the implementation imperative is therefore stronger in organizations that do engage in top-down exploration, and that if exploration is equally costly at the top and the bottom, then top-down exploration (and a strong implementation imperative) is to be preferred (because optimal $\tau_M = 0.2$ when $\tau_S \sim 0$, whereas optimal $\tau_S = 0.5$ when $\tau_M \sim 0$, assuming no observation error or miscommunication: there are several conceivable reasons for this e.g. the communication of top down strategy will nullify the sub's knowledge when there is communication error \lambda=0 therefore, the employee will less suffer from the liability of prior experiences of learning). Further our results about the effects of observation error and miscommunication are qualitatively robust to the introduction of top-down exploration in addition to bottom up exploration.

Robustness Checks

We examined the robustness of our results in two different ways.

First, we looked at the joint effects of all three mechanisms- imperfect communication, observation and top-down exploration, on the existence of the implementation imperative. This is a way of checking whether our understanding of individual mechanisms helps us to make sense of these interactions. To begin with, we look at the interplay of communication and observation error, holding top down exploration to zero. **Figure 5a** shows the consequences of the separation of beliefs and actions when neither M nor S explores. Clearly, mis-communication and observation errors can serve as a valuable source of exploration, when there is no other source of exploration in the organization. As we saw in the results in Figure 3b, mis-communication by itself does not create any exploration benefits; however, it can leverage the effects of exploration. Observation error, on the other hand does create exploration benefits on its own, and these are leveraged by miscommunication. Finally, at high levels of miscommunication, since S effectively ignores M, observation error does not matter. In sum, optimal combinations of

observation error and miscommunication, in organizations with no other source of exploration, involve non-zero values of both.

Next, we consider the case of interaction between observation error and top-down exploration (no communication error). **Figure 5b** shows the same general tendencies as shown in Figure 2a-increasing observation errors worsen performance and create an implementation imperative- i.e. an incentive to improve implementation, even if we allow for some degree of top-down exploration. Finally, we consider the interaction between top down exploration and communication error (no observation error). **Figure 5c** replicates the insight from Figure 3a; as the level of mis-communication increases, there is again an implementation imperative. Additionally, as the level of miscommunication increases, top down exploration, quite naturally becomes irrelevant.

Second, we checked the robustness of our results to different distributional assumptions and the magnitude of noise in the system. We used the normal distribution instead of beta distribution to generate the payoff functions. Our results are qualitatively unchanged. We also checked the impact of increasing noise. Specifically, we increased the level of uncertainty in the bandit by changing the standard deviation of the white noise term (originally e N(0,1)) from s.d.=1 to 1.2 and 1.5 and looked at the cumulative performance with the increasing bottom up exploration across different levels of observation errors. The results show that that upto a certain levels of noise (e.g. s.d. =1.2) the overall patterns are robust and consistent but eventually the organization is overwhelmed if the environmental uncertainty reaches beyond a certain level. For instance, if s.d. is set as much as 1.5, it is found that the organizational learning is no effective anymore and both the manager and subordinate cannot learn from their choices of actions. Thus, for the task environment to be sufficiently noise free to allow organizational learning is a boundary condition

for our arguments. Put differently, learning must be feasible for our central insight - that effective enables the discovery of better strategies by allowing more effective learning from feedback on the value of current strategies- to be relevant.

Conclusions

On the face of it, the glorification of "relentless execution" in a world in which the strategies being executed are rarely infallible appears puzzling. Yet, as our analysis shows, it may be a sensible stance in a world in which beliefs and actions- i.e. strategies and their implementationare separated across organizational actors. While our model is conceptualized in terms of the iterative process of strategy formation and implementation (through organization redesign, for instance- see Galbraith and Kazanjian, 1988), the analytical structure more generally captures any process of experiential learning with separation of beliefs and actions – either because the actions are delegated to somebody other than the individual who holds the beliefs being enacted (e.g. as in strategy implementation) or more generally, because the realized action deviates from intention. In particular, we highlight that none of our results necessarily require that the two agents M and S be hierarchically related; what matters is that beliefs and actions are separated, and while hierarchical ordering is one way in which this is achieved, a division of labour in terms of idea generation and execution among peers would have the same effects.

The key results from our model are that the implementation imperative becomes stronger when a) it is harder for top managers to observe the actions of their subordinates b) when it is harder for top managers to communicate their strategy effectively to their sub-ordinates and c) when top managers also engage in exploratory efforts of their own. The central insight is that when there is a separation of belief and action, as is frequent in organizations because of delegation, effective implementation has benefits beyond the well-known effect of enabling exploitation of good strategies. It also enables the discovery of better strategies by allowing more effective learning from feedback on the value of current strategies.

This occurs through three mechanisms. First, good implementation of currently understood strategy by subordinates avoids the attribution problem that confronts the strategist who cannot perfectly observe if subordinates have indeed executed the desired strategy-whether to ascribe unexpected outcomes to the strategy or to the implementation. Better implementation allows the strategist to escape this inference problem- he can discard poor beliefs and possibly engage in top-down exploration for better strategies. To highlight the intuition of this underlying mechanism, consider the following illustration of the CEO's dilemma when faced with an underperforming business unit:

"Could the unit's lackluster performance have more to do with a mistaken strategy than poor execution? More important, what should he do to get better performance out of the unit? Should he do as the general manager insisted and stay the course-focusing the organization more intensely on execution - or should he encourage the leadership team to investigate new strategy options?" (Mankins and Steele, 2005)

We may think of this as an attribution problem – whether to attribute divergence between expected and actual performance to mistaken strategy or faulty implementation – that is a central feature of experiential learning with difficult to observe actions. Our results show that as this attribution problem becomes more severe as observability of action declines (as μ increases) - it can dwarf any bottom-up exploration benefits of imperfect implementation. This suggests one reason why effective implementation may be useful even with bad strategy, and why the implementation imperative may be so widespread in practice. This is because good implementation solves the attribution problem highlighted in the example above, allowing easier rejection of poor strategies (and possibly a search for better ones), as well as make better use of existing good strategies. These two benefits come at the expense of any potential gains from bottom-up exploration that poor implementation might have allowed.

A corollary of our results is that if an organization wishes to benefit from bottom-up exploration that results from imperfect implementation (i.e. by increasing τ_5), it must also invest in improving its ability to measure implementation success (i.e. by reducing μ). Put differently, if organizations wish to benefit from the bottom-up exploration that arises from imperfect implementation they must nonetheless invest in making the actual actions of employees easier to observe for top managers. This subtle balancing act between autonomy of action but closeness of observation is not an easy one to pull off. Edmondson (2008) draws on her research in hospitals to observe that:

"IHC (a hospital), for example, recognized that physicians, as highly educated experts, might resist process guidelines developed by a committee. For that reason and others, IHC does not discourage doctors from deviating from the guidelines. In fact, the organization invites them to, anytime they judge that good patient care requires it. The only condition: They have to help IHC learn by entering into the computer what they did differently -- and why. This valuable feedback is captured in the system and periodically used by the expert teams to make updates or refinements.....

She goes on to note a little further that:

"It's not easy for a hospital, or any other organization facing cost constraints, to do this. Disciplined reflection takes productive resources off-line, and conventional management wisdom can't help but see this as lost productivity. Nonetheless, the only way to achieve and sustain excellence is for leaders to insist that their organizations invest in the slack time and resources that support this step."

A second advantage of good implementation is that it prevents excessive exploration at the aggregate organizational level because of the inadvertent exploration introduced by the imperfect communication of strategy from senior managers to sub-ordinates. If the downward communication process of strategy is unavoidably noisy, then improving implementation may help to improve organizational performance by reducing what is in effect sub-optimally excessive exploration. While mis-communication and observation error may seem like two sides of the same coin (as both result from the separation of beliefs and actions), there are fundamental distinctions between the two. Imperfect communication leverages bottom-up exploration, but does not contribute any exploration by itself- it merely undampens bottom-up exploration that would have resulted from perfect communication of strategy. Observation error on the other hand is itself a significant source of exploration for the organization. However, while the effect of mis-communication on the knowledge of the strategist is benign, observation errors introduce biases. The implications is that a strategist who is a great communicator but cannot observe the quality of implementation is more harmful to the organization, than one who is a poor communicator but a good observer of implementation performance.

Third, good implementation also prevents excessive exploration at the organizational level when top managers are also exploring- and indeed this is particularly important when there are observation errors and mis-communication of strategy. Under these circumstances, there is a strong implementation imperative. Indeed we can say that for the same cost of bottom-up and top-down exploration, top-down exploration (and a strong implementation imperative) is to be preferred than vice versa. This is primarily because the communication of strategy is effectively a dampener on bottom-up exploration.

These benefits of implementation as a stimulant to organizational adaptation suggest a rationale for the implementation imperative, in the sense that effective implementation may be useful even with bad initial strategies propounded by the senior managers for each of these reasons. Viewing strategy implementation as a learning process for the strategist generates a fundamentally different insight about the value of implementation: unlike the static case when the intuition says: "good beliefs, good implementation; bad beliefs, bad implementation", with learning by the strategist, good implementation can be useful even with bad beliefs.

Our results point to two broader implications for theory and practice. First, given the separation of belief and action (leading to some degree of communication and observation errors between the senior managers and their subordinates), top down exploration enjoy some natural performance advantages over bottom-up exploration. This is quite unrelated to any dominance, status or firm-wide coordination related advantages that top managers enjoy but is rather the simple consequence of the attribution problem that arises because of the separation of beliefs and actions. Our results indicate that a top-down approach favoring the perfect implementation of strategy as designed by top-managers may typically dominate the bottom-up exploration for new strategies in terms of organizational performance. This is because while bottom-up exploration through imperfect implementation may produce good ideas, they are unlikely to be able to be implemented effectively, not is their value likely to be discovered easily. Note that our results suggest that the implementation imperative would be strong even in a world in which there is zero top-down exploration.

In sum, the theory we develop explains why the relentless pursuit of effective implementation may be useful even in a world in which the strategies being implemented are far from optimal. The rationale is tied closely to the relaxation of the unitary actor assumption to account for the fact that actions and beliefs are often separated in organizations. Our analysis shows that this separation between beliefs and actions can increase the value of effective implementation even when the top-down strategies to be implemented are of poor quality and there are potential gains from bottom-up exploration through deviating from the intended strategy.

References

Alexander, Larry D. "Successfully implementing strategic decisions." Long Range Planning 18.3 (1985): 91-97.

Amburgey, T. L., Kelly, D., & Barnett, W. P. (1993). Resetting The Clock: The Dynamics of Organizational Change and Failure. *Administrative Science Quarterly*, *38*(1), 51-73

Augsdorfer, Peter. "Bootlegging and path dependency." Research Policy 34.1 (2005): 1-11.

Axelrod, R., & Cohen, M. D. (2000). *Harnessing complexity: organizational implications of a scientific frontier*. New York: Basic Books.

Baker, G. (2002). Distortion and Risk in Optimal Incentive Contracts. *The Journal of Human Resources*, 37(4), 728-751.

Barnard, C. (1938). The Functions of the Executive. Harvard University Press.

Benner, M. J., & Tushman, M. (2002). Process Management and Technological Innovation: A Longitudinal Study of the Photography and Paint Industries. *Administrative Science Quarterly*, 47(4), 676-706. Johnson Graduate School of Management, Cornell University.

Benner, M. J., & Tushman, M. L. (2003). Exploitation, Exploration, and Process Management: The Productivity Dilemma Revisited. *The Academy of Management Review*, 28(2), 238-256.

Bordley, R. F., & Kirkwood, C. W. (2004). Multiattribute Preference Analysis with Performance Targets. *Operations Research*, 52(6), 823-835.

Breene, R Timothy S, Paul F Nunes, and Walter E Shill. "The chief strategy officer." Harvard business review 85.10 (2007): 84.

Camerer, Colin, and Teck Hua Ho. "Experience-weighted Attraction Learning in Normal Form Games." Econometrica 67.4 (2003): 827-874.

Chandler, A. D. (1962). Strategy and structure.MIT press.

Cyert, R., & March, J. (1963). Behavioral Theory of the Firm (Vol. 2). Malden, MA: Wiley-Blackwell.

Daw, Nathaniel Daw, Nathaniel D., et al. "Cortical substrates for exploratory decisions in humans." Nature 441.7095 (2006): 876-879.D., et al. "Cortical substrates for exploratory decisions in humans." Nature 441.7095 (2006): 876-879.

Denrell, J., & March, J. G. (2001). Adaptation as Information Restriction: The Hot Stove Effect. *Organization Science*, *12*(5), 523-538.

Domjan, Michael, and James W. Grau. The principles of learning and behavior. Wadsworth Publishing Company, 2009.

Edmondson, Amy C. "The competitive imperative of learning." Harvard business review 86.7/8 (2008): 60.

Eisenhardt, Kathleen M., and Shona L. Brown. "Patching. Restitching business portfolios in dynamic markets." Harvard business review 77.3 (1999): 72.

Galbraith, Jay R. Designing complex organizations. Addison-Wesley Longman Publishing Co., Inc., 1973.

Galbraith, Jay R., and Robert K. Kazanjian. "Strategy, technology, and emerging organizational forms." Futures of organizations: Innovating to adapt strategy and human resources to rapid technological change 29 (1988): 41.

Gavetti, G., & Levinthal, Daniel. (2000). Looking Forward and Looking Backward: Cognitive and Experiential Search (Vol. 45, pp. 113-137).

Gavetti, Giovanni, et al. "The Behavioral Theory of the Firm: Assessment and Prospects." The Academy of Management Annals 6.1 (2012): 1-40.

Gavetti, Giovanni. "PERSPECTIVE—Toward a Behavioral Theory of Strategy." Organization Science 23.1 (2012): 267-285.

Gittins, J. C. (1979). Bandit Processes and Dynamic Allocation Indices. *Journal of the Royal Statistical Society*. *Series B (Methodological)*, *41*(2), 148-177. Blackwell Publishing for the Royal Statistical Society.

Gittins, J. C. (1981). Multiserver Scheduling of Jobs with Increasing Completion Rates. Journal of Applied Probability, 18(1), 321-324.

Govindarajan, Vijay. "A contingency approach to strategy implementation at the business-unit level: Integrating administrative mechanisms with strategy." Academy of management Journal 31.4 (1988): 828-853.

Greve, H. R. (2003). A Behavioral Theory of R&D Expenditures and Innovations: Evidence from Shipbuilding. *Academy of Management Journal*, 46(6), 685-702.

Gulati, R., & Puranam, P. (2009). Renewal Through Reorganization: The Value of Inconsistencies Between Formal and Informal Organization. *Organization Science*, 20(2), 422-440.

Guth, William D., and Ian C. MacMillan. "Strategy implementation versus middle management self-interest." Strategic Management Journal 7.4 (2006): 313-327.

Herriott, S. R., Levinthal, Daniel, & March, J. G. (1985).Learning from Experience in Organizations. *The American Economic Review*, 75(2), 298-302.

Hrebiniak, L. G. (2006). Obstacles to effective strategy implementation. *Organizational Dynamics*, 35(1), 12–31. Elsevier.

Johnson, L. K. (2004). Breaking Up Corporate Fiefdoms. Harvard Management Update, 9(12), 3-4.

Kaplan, R. S., & Norton, D. P. (2005). The office of strategy management. Harvard Business Review, 83(10), 72-80.

Kauffman, S. (1993). *The Origins of order: self-organization and selection in evolution*. New York;; Oxford: Oxford University Press.

Kerr, S. (1975).On the Folly of Rewarding A, While Hoping for B.*The Academy of Management Journal*, 18(4), 769-783.

Kim, W. Chan, and Renee A. Mauborgne. "Procedural justice, attitudes, and subsidiary top management compliance with multinationals' corporate strategic decisions." Academy of management journal (1993): 502-526.

Knight, Kenneth E. "A descriptive model of the intra-firm innovation process." The Journal of Business 40.4 (1967): 478-496.

Lave, C., & March, J. (1993). An introduction to models in the social sciences. Lanham Md.: University Press of America.

Lawrence, Paul R., and Jay W. Lorsch. "Differentiation and integration in complex organizations." Administrative science quarterly (1967): 1-47.

Levinthal, D. (1988). A survey of agency models of organizations• 1. *Journal of Economic Behavior& Organization*, 9(2), 153–185.

Levinthal, Daniel A., and James G. March. "The myopia of learning." Strategic management journal 14.S2 (1993): 95-112.

Levinthal, Daniel, & March, J. G. (1981). A model of adaptive organizational search. *Journal of Economic Behavior& Organization*, 2(4), 307-333.

Maier, S. F., & Seligman, M. E. (1976).Learned helplessness: Theory and evidence. Journal of Experimental Psychology: General, 105(1), 3-46.

March, J. G. (1991). Exploration and Exploitation in Organizational Learning. Organization Science, 2(1), 71-87.

March, James G., and Herbert Alexander Simon. "Organizations." (1958).

Merton, Robert K., and Elinor Barber. The travels and adventures of serendipity: A study in sociological semantics and the sociology of science. Princeton University Press, 2003.

Milgrom, P., & Roberts, J. (1990). The Economics of Modern Manufacturing: Technology, Strategy, and Organization. *The American Economic Review*, 80(3), 511-528.

Miner, A. S. (1990). Structural Evolution Through Idiosyncratic Jobs: The Potential for Unplanned Learning. *Organization Science*, 1(2), 195-210.

Mintzberg, H. (1979). The structuring of organizations: a synthesis of the research. Englewood Cliffs N.J.: Prentice-Hall.

Nadler, David, and Michael Tushman. Competing by design: The power of organizational architecture. Oxford University Press, USA, 1997.

Neilson, G. L., Martin, K. L., & Powers, E. (2008). The Secrets to Successful STRATEGY EXECUTION. (cover story). *Harvard Business Review*, 86(6), 60-70.

Nelson, R., &Winter, S. G. (1994). An evolutionary theory of economic change. Cambridge Mass.: Belknap Press of Harvard University Press.

Nickerson, J. A., & Zenger, T. R. (2002). Being Efficiently Fickle: A Dynamic Theory of Organizational Choice. *Organization Science*, 13(5), 547-566.

Noble, C. H., & Mokwa, M. P. (1999). Implementing Marketing Strategies: Developing and Testing a Managerial Theory. *The Journal of Marketing*, 63(4), 57-73.

Nonaka, Ikujiro, and Hirotaka Takeuchi. The knowledge-creating company: How Japanese companies create the dynamics of innovation. Oxford University Press, USA, 1995.

Nowak, Martin, and Karl Sigmund. "A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game." Nature 364.6432 (1993): 56-58.

Oxman, J. A., & Smith, B. D. (2003). The Limits of Structural Change. MIT Sloan Management Review, 45(1), 77-82.

Posen, Hart E., and Daniel A. Levinthal. "Chasing a moving target: Exploitation and exploration in dynamic environments." Management Science 58.3 (2012): 587-601.

Sastry, M. A. (1997). Problems and Paradoxes in a Model of Punctuated Organizational Change. Administrative Science Quarterly, 42(2), 237-275.

Schendel, D., & Hofer, C. (1979). Strategic Management: A New View of Business Policy and Planning. Little Brown.

Simon, H. (1947). Administrative behaviour; 3rd edition, 1976.

Simon, Herbert A. "Models of man; social and rational." (1957).

Simon. (1951). A Formal Theory of the Employment Relationship. Econometrica, 19(3), 293-305.

Skinner, W. (1986). The Productivity Pradox. Management Review, 75(9), 41.

Slater, S. F., Hult, G. T. M., & Olson, E. M. (2010). Factors influencing the relative importance of marketing strategy creativity and marketing strategy implementation effectiveness. *Industrial Marketing Management*, 39(4), 551-559.

Sull, D. N., &Spinosa, C. (2007). Promise-based management: the essence of execution. *Harvard business review*, 85(4), 78-86, 141.

Sutton, R., &Barto, A. G. (1998). Reinforcement learning: an introduction. Cambridge Mass.: MIT Press.

Thorndike, E. L. (1998). Animal intelligence: An experimental study of the associate processes in animals. *American Psychologist*, 53(10), 1125-1127.

Tushman, Michael L., and Charles O'Reilly. "Research and relevance: Implications of Pasteur's quadrant for doctoral programs and faculty development." Academy of Management Journal 50.4 (2007): 769-774.

Yang, Li, Sun Guohui, and Martin J. Eppler. "Making strategy work: A literature review on the factors influencing strategy implementation." Handbook of research on Strategy Process (2009): 165-181.

Zajac, E. J., Kraatz, M. S., &Bresser, R. K. F. (2000). Modeling the dynamics of strategic fit: A normative approach to strategic change. *Strategic Management Journal*, 21(4), 429.

Key Model parameters	Name	Description
T _M	Top-down exploration by Manager	Increasing value indicates greater likelihood that M selects a strategy other than the most attractive one in M's current representation
λ	Imperfect Communication of strategy by Manager to Subordinate	Increasing value indicates that M's selected strategy is NOT taken into account by S.
μ	Manager's probability of making an error when observing Subordinates actions	Increasing value indicates greater likelihood that M believes that S has implemented one of the n-1 strategies other than the one S actually implemented
$1/\tau_S$	Implementation competence	Decreasing value indicates greater likelihood that S selects a strategy other than the most attractive one in S's current representation. Bottom up exploration is τ_S

Table 1: Key parameters (See also Figure 1)

Figure 1 Model structure



Manager who explores for new strategy with exploration parameter τ_{M}

3

Figure 2a. Cumulative organizational performance ($\Sigma \pi$) and bottom-up exploration τ_s with varying observation errors (μ) and ($\lambda =0, \tau_M=0$)



Figure 2b. Propensity to explore and Knowledge level and bottom-up exploration τ_s with observation errors (μ =0), and (λ =0, τ_M =0)



Figure 2c: Propensity to explore and Knowledge level with varying observation errors (μ), and ($\lambda = 0$, $\tau_M = 0$, $\tau_S = 0.01$ or 0.5)



Figure 2d. The substitutive relationship between bottom-up exploration (τ_s) and observation errors (μ)



Figure 3a. Cumulative organizational performance ($\Sigma \pi$) and bottom-up exploration τ_s with imperfect communication (λ), and ($\mu = 0, \tau_M = 0$)



Figure 3b. Propensity to explore and Knowledge level with varying degrees of imperfect communication (λ), and ($\mu = 0$, $\tau_M = 0$, $\tau_S = 0.01$ or 0.5)



Figure 3c. The relationship between optimal bottom-up exploration (τ_s) and imperfect communication (λ), ($\mu = 0, \tau_M = 0$)



Figure 4a Cumulative organizational performance ($\Sigma \pi$) and bottom-up exploration τ_s with top down exploration ($\tau_M > 0$) and no mis-communication ($\lambda = 0$), or observation error ($\mu = 0$)



Figure 4b Propensity to explore and Knowledge level with varying degrees of bottom up exploration (τ_s), and ($\mu = 0, \tau_s = 0.01$ or 0.5, $\lambda = 0.0$)



Figure 4c. Cumulative average performance over the varying top down (τ_M) and bottom up exploration (τ_s) with no observation errors ($\mu = 0$) and no mis-communication ($\lambda = 0.0$).



Figure 5a. Cumulative average performance over the varying misalignment(λ) and observation error(μ) when both bottom exploration (τ_s) and top down (τ_m) are low ($\tau_M = \tau_s = 0.01$).



Figure 5b: Cumulative average performance over the varying top down (τ_M) and bottom up exploration (τ_S) with varying observation errors (0<= μ <=0.3) and no mis-communication (λ =0.0).





Figure 5c Cumulative average performance over the varying top down (τ_M) and bottom up exploration (τ_s) across different level of mis-communication ($0 \le \lambda \le 1$). The degree of observation error is kept constant ($\mu = 0$).





Europe Campus Boulevard de Constance 77305 Fontainebleau Cedex, France Tel: +33 (0)1 60 72 40 00 Fax: +33 (0)1 60 74 55 00/01

Asia Campus 1 Ayer Rajah Avenue, Singapore 138676 Tel: +65 67 99 53 88 Fax: +65 67 99 53 99

Abu Dhabi Campus Muroor Road - Street No 4 P.O. Box 48049 Abu Dhabi, United Arab Emirates Tel: +971 2 651 5200 Fax: +971 2 443 9461

www.insead.edu



The Business School for the World®