Paper to be presented at the

35th DRUID Celebration Conference 2013, Barcelona, Spain, June 17-19

Under-Confidence: Second-Order Knowledge and the Efficacy of Learning

Hart Posen
University of Wisconsin
School of Business
hposen@bus.wisc.edu

Dirk Martignoni
University of Zurich
Department of Business Administration
dirk.martignoni@business.uzh.ch

Markus Lang
University of Zurich
Department of Business Administration
markus.lang@business.uzh.ch

Abstract

Knowledge is a central construct in the modern strategy literature because it contributes to performance differences across firms. In the Carnegie School tradition, knowledge is the outcome of a process of search and learning by which firms discover better solutions to the challenges they face. Knowledge can be interpreted as more accurate beliefs about the merits of alternative policy choices - which we term first-order knowledge. Such knowledge may be employed to allocate resources across alternative investment opportunities, for example, new plants, new drug molecules, or venture capital investments. Yet knowledge has a second dimension related to a firm's confidence in its beliefs. Intuition suggests that accuracy has limited value without confidence, and confidence has limited value without accuracy. We term accuracy-confidence matches second-order knowledge. We identify conditions under which low second-order knowledge (a confidence mismatch - high confidence in inaccurate beliefs) may enhance rather than diminish firm performance because it increases the efficacy of organizational learning. Thus, while knowledge is usually thought of as a unidimensional construct, in fact it embodies two distinct dimensions that interact in non-trivial ways in the organizational learning process: the accuracy of beliefs about the relative merits of alternatives and the confidence in those beliefs.

Jelcodes:M00,-

Under-Confidence:

Second-Order Knowledge and the Efficacy of Learning

Abstract

Knowledge is a central construct in the modern strategy literature because it contributes to performance differences across firms. In the Carnegie School tradition, knowledge is the outcome of a process of search and learning by which firms discover better solutions to the challenges they face. Knowledge can be interpreted as more accurate beliefs about the merits of alternative policy choices — which we term first-order knowledge. Such knowledge may be employed to allocate resources across alternative investment opportunities, for example, new plants, new drug molecules, or venture capital investments. Yet knowledge has a second dimension related to a firm's confidence in its beliefs. Intuition suggests that accuracy has limited value without confidence, and confidence has limited value without accuracy. We term accuracy-confidence matches second-order knowledge. We identify conditions under which low second-order knowledge (a confidence mismatch — high confidence in inaccurate beliefs) may enhance rather than diminish firm performance because it increases the efficacy of organizational learning. Thus, while knowledge is usually thought of as a unidimensional construct, in fact it embodies two distinct dimensions that interact in non-trivial ways in the organizational learning process: the accuracy of beliefs about the relative merits of alternatives and the confidence in those beliefs.

To know that we know what we know and that we do not know what we do to not know, that is true knowledge

— Confucius

1. Introduction

An important strategic challenge facing organizations is the need to make repeated choices across a set of policy alternatives. For instance, firms must allocate resources across alternative investment opportunities and make subsequent milestone-based go/no-go decisions, entrepreneurs must select and then repeatedly revise their strategies, a biotech company must evaluate a molecule and recommit to invest at subsequent rounds of clinical trials, and venture capital firms must repeatedly fund portfolio companies across several stages of development. The efficacy of a firm's choice is an increasing function of the quality of its knowledge – more accurate beliefs about the relative merits of the alternatives should translate into better choices (Simon 1955, March 1991, Levinthal 1997). We term the accuracy of beliefs as first-order knowledge. Yet this simple decision process is confounded by the observation that most repeated choices must be made under uncertainty that makes evaluation of prior outcomes difficult (Knudsen and Levinthal 2007) and reduces the efficacy of learning about the alternatives (March, Sproull, and Tamuz 1991, Levinthal and March 1993, Posen and Levinthal 2012). In such situations knowledge has a second equally important but underexplored component: a firm's confidence in its beliefs (March 2010, March 2011). Confidence is a central component of knowledge because it alters the way in which firms interpret performance feedback from their choices. Intuition suggests that accuracy has limited value without confidence, just as confidence has limited value without accuracy. We call these accuracy-confidence matches, such as high confidence in accurate beliefs, "second order knowledge." We argue that there are conditions under which *low* second-order knowledge (high confidence in inaccurate beliefs) may enhance rather than diminish the efficacy of this type of strategic choice.

The importance of confidence as a dimension of knowledge is inextricably tied to uncertainty in the task environment. In the absence of uncertainty, a "sample of one" on any alternative is sufficient to generate knowledge that is high in both accuracy and confidence (March et al. 1991). The behavioral and cognitive challenges of learning (March and Olsen 1975; Levinthal and March 1993) are magnified when the true value of an alternative cannot be known with certainty (e.g., Denrell 2007; Denrell and March 2001; Knudsen and Levinthal 2007; March 1996; March, Sproull, and Tamuz 1991; Posen and Levinthal 2012). Uncertainty necessitates repeated sampling, through which a firm improves the accuracy of its beliefs as it takes action and receives feedback, but also enhances its confidence in its beliefs.

To clarify the notion of confidence, consider a firm choosing between two alternative technologies, the true value of which is unknown to them. The firm believes that an investment in technology A will generate a profit of \$100 ± \$5 and technology B will generate a profit of \$50 ± \$20. Further assume that the true expected profit from technology A is \$105 and technology B is \$25. Accuracy, first-order knowledge, reflects the extent to which the firm's mean beliefs (\$50, \$100) about these technologies correspond with reality — we refer to this as first-order knowledge. This firm has better first-order knowledge about technology A than it does about technology B. Relatedly, confidence reflects the firm's estimate of the precision of its beliefs — the interval within which the mean value of the alternative is expected to lay (\$10, \$40). Thus, the firm has higher confidence in its accurate beliefs about technology A than it does in its inaccurate beliefs about technology B. This reflects an accuracy-confidence match — high second-order knowledge.

While the management literature has tended to directly address issues of knowledge as the accuracy of beliefs (e.g., Coff 1999; Winter, Denrell, and Fang 2003; Makadok and Barney 2001), issues of confidence have been discussed in a less direct fashion. The literature tends to consider as knowledge the case of high accuracy and high confidence. For example, studies of the value of pre-entry knowledge in an entrepreneurial context (Chen, Williams, and Agarwal 2012; Dencker, Gruber, and Shah 2009; Gruber, MacMillan, and Thompson 2008) implicitly assume that the beliefs that underlay pre-entry knowledge are accurate (otherwise these beliefs would not be performance enhancing), and that the entrepreneurs have high confidence in these beliefs (otherwise they would not act on them).

In studies of knowledge as "the capacity to act" (Weick 1999, Levinthal and Rerup 2006), the accuracy of beliefs takes a backseat to confidence. Consider Weick's (1999) illustration in which Hungarian soldiers trapped in the Pyrenees find there way home using a map of the Alps — they were able to take action although their knowledge was inaccurate (and their happy ending is purely the result of chance). Confidence that the map was accurate is the focal construct, rather than accuracy. Weick's illustration points to the motivational quality of confidence, but not to any particular performance outcome (good or bad). That is, the soldiers may not have marched at all in the absence of confidence in the map. Unanswered is the question of whether low accuracy combined with high confidence (low second-order knowledge) can reliably generate favorable results.

Putting aside anecdotes, systematic research on the costs or benefits of low second-order knowledge (mismatch between accuracy and confidence) is relatively sparse. There are two key reasons for this lack of research. The simplest reason is that the benefits of high second-order knowledge (and by implication, the costs of low second-order knowledge) are taken-for-granted — accuracy-confidence matches are good (and mismatches are bad). In the academic literature, for example, Nonaka and von Krogh (2009) imply that confidence in accurate beliefs is an important "success factor." In the popular management press, for example, Kanter (2006 p.71) focuses on the "expectation of success" where "confidence is based on reasonable expectations."

There is also a theory-based reason for sparse research on low second-order knowledge. In the context of one-shot decision-making, confidence is of little relevance to choice. For a risk neutral firm, confidence may motivate action, but the choice itself is based on beliefs independent of confidence. Beliefs that are inaccurate will, on average, engender uninformed action and poor outcomes, while accurate beliefs will engender superior outcomes. In a one-shot decision, confidence only binds on choice if the firm is risk averse, where confidence affects choice in a straightforward manner.

Yet many important empirical contexts involve opportunities for learning across repeated decisions. Consider knowledge employed to allocate resources across alternative strategic investment opportunities, for example, new plants, new drug molecules, or venture capital investments. These strategic choices reflect the need to make repeated decisions across time, where intermediate feedback embodies significant uncertainty. After each decision, the firm receives noisy feedback from the world (revenue, profits, productivity, defects, etc.) that is used to update its beliefs about the relative merits of alternatives. In turn, these more accurate beliefs are the basis of subsequent choice. Our central assumption, following March (2010), is that processes of learning alter both the accuracy and confidence of beliefs. When firms learn under conditions of uncertainty, the coevolution of accuracy and confidence gives rise to the possibility that low second-order knowledge may sometimes be performance enhancing.

To understand the implications of second-order knowledge in this context, we anchor our formal development on the multi-armed bandit model. This model is the canonical representation of learning under conditions of uncertainty (Denrell and March 2001; Holland 1975; March 1996, 2003, 2010; Posen and Levinthal 2012). The bandit model takes its name from a slot machine analogy in which a firm seeks to maximize the flow of returns over time. In each period, the firm chooses one alternative from a set of policy alternatives where the payoff to an alternative is drawn from a given probability distribution. The firm does not know the payoffs to the alternatives but can form beliefs about them, which are refined over time through a process of learning. Based on its beliefs about the relative merits of alternatives, the firm makes a choice from the set of alternatives, receives feedback from the environment in the form of an outcome signal and updates its beliefs. Thus, the firm is portrayed as possessing a mental model or cognitive map (derived from its own prior experience) that encapsulates its understanding of the merits of the available set of choices. We build on this canonical model by formally considering both the firms mean beliefs about the relative merits of alternatives, but also the confidence associated with those beliefs. We endow firms with different initial levels of accuracy and confidence and examine the implications for learning.

The intuition emerging from our model reflects the observation that in the process of learning under uncertainty, confidence in beliefs is central to moderating the extent to which a firm learns from noisy feedback and thus the efficacy of allocating resources across investment opportunities. Uncertainty has two effects on learning — it endogenously increases the extent to which a firm explores (Denrell and March 2001), and at the same time reduces the efficacy of evaluating feedback (Knudsen and Levinthal 2007). Confidence alters the relative balance of these two effects, and in doing so, moderates the quality of learning. In particular, under high levels of uncertainty, high confidence reduces the extent to which a firm responds, both in terms of action (exploring other alternatives) and evaluation (listing to, and incorporating, feedback). Thus, in contrast to Confucian wisdom on second-order knowledge, that we must know that "we do not know what we do not know," sometimes it is valuable not to know that we do not know.

This paper proceeds as follows. In the next section, we describe the theoretical background and setup of our multi-armed bandit computational model. In section 3, we present the results of simulation experiments in which we examine the properties of the model of experiential learning as a function of initial belief accuracy and confidence. In the final section, we discuss the implications of these results for theory and practice.

2. Model

To examine the implications of second-order knowledge, we implement a standard multi-armed bandit model. The bandit model (Gittins 1979, Berry and Fristedt 1985) has been the subject of significant study because its underlying structure closely resembles a variety of realistic economic situations ranging from research settings such as R&D (Hardwick and Stout 1992), to strategic issues, such as product pricing (Bergemann and Välimäki 1996), and consumer choice (Gans, Knox, and Croson 2007). There are two common features underlying economic problems that are modeled in a bandit framework. First, information about the returns to an alternative can only be gathered by sampling it. Second, feedback from trials is subject to uncertainty that gives rise to variation in possible outcomes, and as such, any particular experience is likely to be misleading.

Formally, the bandit model reflects a sequential choice problem where, at each point in time, t, a firm must choose among N alternatives. The payoffs to alternative $i=\{1,\ldots,N\}$ is drawn from a standard normal distribution such that the state of the environment can be described by the vector of payoffs to the alternatives, $P=[p_1,\ldots,p_N]$ where $p_i \sim N(0,1)$. The feedback (reward signal) r_i from alternative i is normally distributed with $r_i \sim N(p_i, \eta^2)$. That is, the mean of the signal r_i equals the actual payoff of alternative i and the variance of the signal is $\eta^2>0$. We interpret the variance η^2 in the signal as the uncertainty in the environment where a higher value of η^2 corresponds to a more uncertain environment.

The firm does not know the payoffs to the alternatives but can form beliefs about them. The firm's beliefs, $b_{i,t}$, about alternative i at time t are encapsulated by the vector $B_t=[b_{1,t},\ldots,b_{N,t}]$. To refine its beliefs and maximize the value of the stream of payoffs, the firm engages in learning.

In the initial period t=0, the firm has initial beliefs, $B_0=[b_{1,0},\ldots, b_{N,0}]$, about the payoffs of the N alternatives $P=[p_1,\ldots, p_N]$. We differentiate between two polar cases regarding these initial beliefs about the merits of each alternative: (i) The firm has accurate initial beliefs across the alternatives in the sense that its beliefs match the payoffs of the alternatives, i.e., $B_0 = P$ (high accuracy). (ii) The firm has non-informative or random initial beliefs across the alternatives in the sense that its beliefs are drawn from the same distribution as the payoffs, i.e., $B_0=[b_{1,0},\ldots, b_{N,0}]$ where $b_{i,t} \sim N(0,1)$ for alternative i (low accuracy).

In each subsequent period, the firm makes a choice from the set of alternatives and receives feedback from the environment in the form of an outcome signal. For the sake of simplicity, we assume that the firm always chooses the alternative that he believes is associated with the highest payoff ("greedy" choice rule). Our results are qualitatively unchanged when a less greedy choice rule is employed.

Consistent with stochastic learning models (Bush and Mosteller 1955, Lave and March 1993, Sutton and Barto 1998, Denrell and March 2001), we assume that the firm's beliefs $b_{i,t}$ about the merits of alternative i at time t is proportional to the difference between $b_{i,t-1}$ and the signal $r_{i,t}$. Specifically, the firm's beliefs are updated according to

$$b_{i,t} = b_{i,t-1} + \alpha\,(r_{i,t} - b_{i,t-1}), \tag{1}$$

where $0 \leq \alpha \leq 1$ is a weighting parameter. From this, we can derive the distribution of the beliefs about the merits of alternative i at time t as

$$b_{i,t} \sim N(\mu_{i,t}\,,\, \sigma_t^2) \quad \text{with} \quad \mu_{i,t} = (1-\alpha)^{ki}\, b_{0,i} + \alpha^{ki}\, p_i \;\; \text{and} \;\; \sigma_t^2 = [1-(1-\alpha)^{ki}]^2\, \eta^2, \tag{2}$$

where $k_i > 0$ is the number of trials associated with alternative i until period t. Clearly, in the limit, the beliefs about the merits of alternative i converge to a normal distribution with mean p_i and variance η^2, i.e., $\lim_{k \to \infty} b_{i,t} = b_i \sim N(p_i\,,\, \eta^2)$. Thus, we can interpret the standard deviation σ_t of the distribution in Equation (2) as a measure for the level of confidence that the firm holds in its beliefs. As a consequence, we can express Equation (1) in terms of confidence as

$$b_{i,t} = \frac{\sigma_t}{\eta} b_{0,i} + \left(1 - \frac{\sigma_t}{\eta}\right) r_{i,t}. \tag{3}$$

A large (small) parameter σ_t represents a firm that has a high (low) level of confidence in its beliefs. In the subsequent analysis, we write $\gamma = \frac{\sigma_1}{\eta}$ and refer to $0 \leq \gamma \leq 1$ as the firms normalized level of initial confidence.

3. Analysis

In the following sections, results are reported for the case of 20 alternatives with their payoffs drawn from a standard normal distribution. Each experiment involves 50,000 replications. If not mentioned otherwise, we report long-run performance in t=500. In our analysis, we examine the interplay between the accuracy of beliefs and the confidence in these beliefs across different levels of environmental uncertainty.

Interplay of Confidence and Uncertainty

In Figures 1 and 2, we seek to understand the basic properties of the interplay between first-order and second-order knowledge. While Figure 1 shows the cross-sectional results for all levels of uncertainty, Figure 2 reports the time series for some selected values of uncertainty and confidence for firms with inaccurate beliefs. The left panel in Figure 1 reports the average long-run performance in t=500 (y-axis) for firms endowed with highly accurate initial beliefs about the payoffs of the alternatives they are facing as a function of the level of uncertainty (x-axis), ranging from zero or absence of uncertainty ($\eta^2 = 0$) to high uncertainty ($\eta^2 = 5$). We differentiate between firms that have a low level of confidence (γ=0.1, dotted line) and a high level of confidence (γ=0.9, dashed line) in their beliefs. The right panel in Figure 1 is analogous to the left panel; the only difference is that firms are endowed only with low accurate (non-informative or random) initial beliefs about the value of the payoffs.

< Insert Figure 1 about here >

The results for firms with accurate beliefs (reflecting high first-order knowledge) in the left panel are not surprising: for all levels of uncertainty, having high second-order knowledge (high accuracy and high confidence) is always better than having low second-order knowledge (high accuracy and low confidence). Obviously, if firms are endowed with fully accurate beliefs, there is no value in learning. Put differently, if endowed with accurate beliefs, firms should ideally also have high confidence in their beliefs and refrain from updating these beliefs based on noisy feedback. With accurate beliefs, any updating or learning can only have negative implications. There is no upside to learning if beliefs are already accurate. These negative implications are particularly pronounced if firms have to learn from very noisy signals (high uncertainty).

In contrast, for inaccurate initial beliefs (reflecting low first-order knowledge), the interaction between confidence and uncertainty is less obvious (right panel). In environments of low uncertainty, firms are best off having high second-order knowledge (low accuracy and low confidence). In other words, there is a match between accuracy and confidence. Surprisingly, in environments of high uncertainty, a mismatch between accuracy and confidence actually enhances rather than diminishes performance: ideally, firms should have high confidence in inaccurate beliefs (low second-order knowledge).

Given that the counter-intuitive implications of second-order knowledge are in the right panel in Figure 1, we focus our attention in the subsequent analysis on low first-order knowledge (i.e., when initial beliefs are inaccurate).

Figure 2 displays the evolution of average performance (y-axis) over time until t=5,000 (x-axis) for firms with inaccurate beliefs in an environment with a low level of uncertainty ($\eta^2 = 1$, left panel) and a high level of uncertainty ($\eta^2 = 4$, right panel). In both panels, the dotted and dashed lines report the average performance for firms with low confidence (γ=0.1) and high confidence (γ=0.9), respectively.

< Insert Figure 2 about here >

Figure 2 shows that the basic pattern of Figure 1 holds for all time-horizons. For firms with inaccurate beliefs, low confidence is always better than high confidence in environments of low uncertainty, while high confidence is always better than low confidence in environments of high uncertainty. Indeed, the performance difference between firms with low and high confidence increases even over time.

In the subsequent analysis, we seek to explain why inaccurate initial beliefs and high confidence (low second-order knowledge) engender superior performance under high levels of uncertainty. In particular, we will uncover the mechanisms that give rise to the performance implications of a mismatch between accuracy and confidence.

In models of learning under uncertainty, the efficacy of a learning mechanism is a function of two components: alternative generation and alternative evaluation (Denrell and March 2001, Knudsen and Levinthal 2007). First, the ability to generate new alternatives increases in uncertainty and uncertainty induces exploration (e.g., Denrell and March 2001). In particular in greedy systems, which tend to overemphasize exploitation, this induced exploration through uncertainty can have positive performance implications. Second, the presence of uncertainty makes evaluation of alternatives much more prone to errors (Knudsen and Levinthal 2007). Consider a setting in the absence of uncertainty. The feedback from sampling an alternative (the performance signal) matches the actual payoff of the chosen alternative. As a result, there are neither errors of under- or over-estimation nor errors of omission or commission. In the presence of uncertainty, however, the signal-to-noise ratio in the feedback declines. Feedback from sampling a particular alternative may indicate a higher (or lower) payoff than the actual (expected) payoff of this alternative. Assessing the true value of alternatives becomes more difficult and prone to errors in the presence of uncertainty.

In Figure 3, we decompose the performance effect of Figure 1 for inaccurate beliefs into these two components: (1) an exploration effect induced through uncertainty and (2) evaluation effect driven by the increased difficulty in evaluating alternatives under uncertainty. The dashed line reflects the performance implications of the exploration effect, while the dotted line displays the performance implications for the evaluation effect. The exploration effect reflects the average performance of the best alternative a firm has ever sampled in the first 500 periods. We isolate the evaluation effect by identifying the best alternative a firm has ever sampled in the first 500 periods and subtracting its payoff from the payoff of the alternative that the firm is sampling by period 500. Thus, this effect is zero if the best alternative the firm has sampled in the first 500 periods is the one the firm is still sampling in period 500. It becomes negative if the firm has sampled a better alternative but – due to misleading performance signals – has abandoned it again in subsequent periods. The net effect of the evaluation and exploration effects is plotted as the solid line, fully reconstructing the main result in the right panel of Figure 1. We display the normalized performance effect (i.e., average long-run performance less than that achieved in the absence of uncertainty) for different levels of uncertainty (x-axis) for firms with low levels of confidence (γ=0.1, left panel) and high levels of confidence (γ=0.9, right panel).

< Insert Figure 3 about here >

Results suggest that for all levels of uncertainty and confidence, the exploration effect is always positive, while the evaluation effect is always negative. However, their relative contributions change in uncertainty and confidence. In particular, the level of confidence crucially determines the strength of the negative evaluation effect and in turn the normalized performance. For a given level of uncertainty, the negative evaluation effect is stronger for low than for high levels of confidence. This differential effect of the evaluation effect explains why the normalized performance for firms with low confidence turns negative above a certain level of uncertainty. In contrast, for firms with high confidence, the negative evaluation effect is not sufficiently strong so that the normalized performance remains positive for all levels of uncertainty. As a result, firms with low second-order knowledge outperform firms with high second-order knowledge for high levels of uncertainty.

In the following subsections, we seek to uncover the mechanisms underlying the evaluation and exploration effects for inaccurate beliefs (low first-order knowledge) and different levels of second-order knowledge. We begin by further decomposing the evaluation effect, and then further decompose the exploration effect.

Decomposition of the Evaluation Effect

In the analysis above, we find that the evaluation effect (dotted line in Figure 3) is consistently negative and its strength is increasing with a higher level of uncertainty. Moreover, the negative evaluation effect is more pronounced for low levels of confidence than for high levels of confidence.

In Figure 4, we further decompose this evaluation effect into two components: (1) the probability that an evaluation error occurs (dashed line, left y-axis) and (2) the average size of the evaluation error (dotted line, right y-axis). Combined, these two components reconstruct the average evaluation effect (solid line) in Figure 3. We distinguish between firms with low levels of confidence (γ=0.1, left panel) and high levels of confidence (γ=0.9, right panel).

< Insert Figure 4 about here >

Not surprisingly, the probability of committing an evaluation error, the dashed line in Figure 4, is increasing in uncertainty, albeit with a decreasing rate. In the absence of uncertainty, firms do not commit evaluation errors and consequently do not abandon the best alternative. With high uncertainty, any signal can be misleading and firms may completely overturn their belief about the merits of the current alternative. In the context of our model, errors of overestimation do not trigger changes in behavior because firms will stick to their current alternative. Hence, errors of underestimation are a necessary condition for triggering changes in behavior. Firms will only abandon their current alternative if they underestimate their true value and believe that it is below the second-best alternative. Yet, the impact of an increase in evaluation errors through higher uncertainty can be compensated for by an increase in confidence. For a given level of uncertainty, the probability of committing an evaluation error is decreasing with confidence.

The dotted line shows that the average size of evaluation errors is almost linearly decreasing with uncertainty. For a high level of uncertainty, a firm may erroneously come to the belief that a very attractive alternative has a low payoff. With lower uncertainty, firms may abandon good but almost never the best alternative. In addition, the average size of evaluation errors is decreasing with confidence for a given level of uncertainty.

Combined, the probability of committing an evaluation error and the size of the evaluation error fully reconstruct the evaluation effect (solid line). The differential effects of confidence on (i) the probability of committing an evaluation error and (ii) the average size of evaluation errors explains why the negative evaluation error is more pronounced for low than for high levels of confidence.

<u>Decomposition of the Exploration Effect</u>

In the analysis above, we find that across all levels of confidence, the positive exploration effect increases with uncertainty albeit with a decreasing rate. We further decompose this exploration effect into two components: (i) the extent of exploration (Figure 5) and (ii) the performance effect of additional exploration (Figure 6). On the one hand, a higher level of uncertainty increases the extent of exploration (Levinthal 1997, Denrell and March 2001). On the other hand, however, the marginal performance effect (and by implication, the average performance effect) of exploration is decreasing in the number of alternatives explored.

< Insert Figures 5 and 6 about here >

Figure 5 reports the average number of different alternatives (y-axis) that were sampled at least once during the first 500 periods (extent of exploration) as a function of the uncertainty in the environment (x-axis) for different levels of confidence. Consistent with Denrell and March (2001), we find that the extent of exploration is increasing in uncertainty for all levels of confidence. This increase, however, is not linear. Instead, the marginal effects are decreasing: increasing uncertainty from 1 to 2 has a larger effect than increasing uncertainty from 4 to 5. We also find that, for a given level of uncertainty, lower confidence leads to broader exploration.

Figure 6 displays the average performance effect (y-axis) as a function of the extent of exploration (x-axis) for different levels of confidence and the whole range of uncertainty (i.e., from 0 to 5). For each level of uncertainty and confidence, we calculate long-run performance and the extent of exploration. From this, we can compute the average performance effect of each explored alternative by dividing the long-run performance through the extent of exploration. The results in Figure 6 provide an intuition for the performance implications of exploration introduced through uncertainty or (lack of) confidence. The average performance effect of each explored alternative is independent of whether this exploration was induced by uncertainty or a lack of confidence. As long as different levels of uncertainty and confidence are associated with similar levels of exploration, the average performance effects are identical.

Figure 6 also offers an intuition about how the endogenous sampling process affects the returns to exploration. With extreme value statistics, we can compute the expected value of sample maxima. In Figure 6, the corresponding values are reflected in the solid line. Consider, for example, the case in which the alternatives' payoffs are drawn from a standard normal distribution. If 5 out of 20 alternatives are sampled, the expected value of the best among these 5 alternatives is 0.25. If just one alternative is sampled, this value drops to the zero, i.e. the mean of the payoff distribution. If the sampling process is endogenous, however, even a sample of just one has positive mean. Consider, for example, that by chance a firm picks the best among twenty alternatives in t=0. If the sampling process is endogenous, a highly confident firm will rarely abandon its choice. If however, the first pick is particularly unattractive, the firm will continue sampling. In more abstract terms, the number of alternatives explored is a function of their quality. In Figure 6, the dotted and dashed lines reflect the effects of the endogenous sampling process and are computed by taking the difference between the solid line (Figure 6) and the dotted and dashed lines (Figure 5) generated through our experiments with the armed bandit model. Following our example above, this difference is largest for small samples (4 alternatives) and decreases to zero for a full sample of 20 alternatives.

Combining Figures 5 and 6, we find that both the returns to exploration and the extent of exploration are increasing in uncertainty, albeit with a decreasing rate. These two mechanisms generate a positive exploration effect for all levels of uncertainty and confidence.

In sum, our model suggests that firms operating in environments with high uncertainty may sometimes benefit from having high confidence in inaccurate beliefs (low second-order knowledge), while firms operating in environments with low uncertainty are better served by having low confidence in inaccurate beliefs (high second-order knowledge). These effects are driven by two interacting mechanisms, the positive exploration effect and the negative evaluation effect. Confidence alters the relative balance of these two effects, and in doing so, moderates the quality of learning. In particular, high confidence reduces both the extent to which a firm explores other alternatives and commits evaluation errors.

4. Discussion

Knowledge is a central construct in the modern strategy literature, contributing to performance differences across firms (Nelson and Winter 1982, Grant 1996). In the Carnegie School tradition, knowledge is the outcome of a process of search and learning by which firms discover better solutions to the challenges they face (Simon 1955; Cyert and March 1963; Nelson and Winter 1982). Knowledge, which in a strategy sense can be interpreted as more accurate beliefs about the merits of alternative policy choices, enables a firm to make more effective strategic choices and as such, engenders superior performance (Simon 1955; March 1991; Levinthal 1997). In this paper, we highlight a second, equally important but underexplored dimension of knowledge: a firm's confidence in its beliefs (March 2010, March 2011).

Intuition suggests that accuracy has limited value without confidence, just as confidence has limited value without accuracy. One might be tempted to conclude that a firm would hope to have high confidence in accurate beliefs and low confidence in inaccurate beliefs. These accuracy-confidence matches reflect second-order knowledge. Our central argument is that second-order knowledge is not always performance enhancing.

Existing research, often implicitly, focuses on instances of high second-order knowledge, i.e. a match between the accuracy of beliefs and confidence in those beliefs. This implicit focus may explain why existing studies often fail to distinguish between optimism (beliefs about the merits of an alternative that are biased positively) and overconfidence (Hmieleski and Baron 2008). Yet, without a conceptual separation between measures of accuracy and measures of confidence, we cannot fully address questions of low-second order knowledge. Our central assumption, following March (2010), is that processes of learning alter both the accuracy and confidence of beliefs. When firms learn under conditions of uncertainty, the coevolution of accuracy and confidence gives rise to the possibility that low second-order knowledge may sometimes be performance enhancing.

The intuition emerging from our model reflects the observation that in the process of learning under uncertainty, confidence in beliefs is central to moderating the extent to which a firm learns from noisy feedback. Uncertainty has two effects on learning — it endogenously increases the extent to which a firm explores (Denrell and March 2001), and at the same time reduces the efficacy of evaluating feedback (Knudsen and Levinthal 2007). Confidence alters the relative balance of these two effects, and in doing so, moderates the quality of learning. As a consequence, there are conditions under which low second-order knowledge (confidence mismatch) is performance enhancing. In contrast to Confucian wisdom on second-order knowledge, that we must know that "we do not know what we do not know," sometimes it is valuable not to know that we do not know.

Our theory is broadly consistent with the behavioral theory of the firm (Cyert and March 1963; March and Olsen 1975) under which boundedly rational firms face challenges in interpreting feedback from the world on the merits of a given policy choice (Greve 2002, 2008; March and Olsen 1975; Hu, Blettner, and Bettis 2011). In contrast to this early work on aspirations, we focus on how confidence alters the process by which feedback is interpreted. In the presence of uncertainty, more often than not, feedback will not (exactly) confirm a firm's beliefs (March 1994). Such "post-decision surprises", sometimes positive and sometimes negative, are an inherent characteristic of choice under uncertainty (Harrison and March 1984). Firms with low confidence in their current beliefs will experience stronger post-decision surprises than firms with high confidence in their current beliefs. Only negative post-decision surprises may result in post-decision regret and a revision of a firm's choice. While in general, post-decision surprises are not necessarily bad, they are bad if a firm's beliefs are already completely accurate. Then, surprises may make a firm abandon a particularly attractive choice. Thus, while Harrison and March (1984) demonstrate that the extent to which firms experience post-decision surprises is a function of the level of uncertainty and number of available alternatives, we point to the moderating effect of confidence and accuracy on the extent to which firms experience post decision surprises.

In sum, knowledge plays a central role in the management literature. We have sought to enhance our understanding of the interplay between first- and second-order knowledge and the implications for the efficacy of learning, and in turn, future performance. Second-order knowledge is central in contexts in which firms make repeated strategic decisions across a set of alternatives. The issues surrounding the implications of knowledge for learning remain a fertile and important line of inquiry for organizational scholars.

Bibliography

Bergemann, D., & Välimäki, J. (1996). Learning and strategic pricing. *Econometrica*, *64*(5), 1125-1149.

Berry, D. A., & Fristedt, B. (1985). *Bandit problems: Sequential allocation of experiments* London: Chapman and Hall.

Bush, R., & Mosteller, F. (1955). Stochastic models for learning. In *Stochastic models for learning* New York: John Wiley & Son.

Chen, P. -L., Williams, C., & Agarwal, R. (2012). Growing pains: Pre-entry experience and the challenge of transition to incumbency. *Strategic Management Journal*, *33*(3), 252-276. doi:10.1002/smj.954

Coff, R. 1999. How buyers cope with uncertainty when acquiring firms in knowledge-intensive industries: Caveat emptor. *Organization Science*, 10(2): 144-161

Cyert, R., & March, J. G. (1963). *A behavioral theory of the firm* Englewood Cliffs, NJ: Prentice-Hall.

Dencker, J., Gruber, M., & Shah, S. (2009). Pre-Entry knowledge, learning, and the survival of new firms. *Organization Science*, *20*(3), 516-537.

Denrell, J. (2007). Adaptive learning and risk taking. *Psychological Review*, *114*(1), 177-187.

Denrell, J., & March, J. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, *12*(5), 523-538.

Denrell, J., Fang, C., & Winter, S. G. (2003). The economics of strategic opportunity. *Strategic Management Journal*, *24*(10), 977.

Gans, N., Knox, G., & Croson, R. (2007). Simple models of discrete choice and their performance in bandit experiments. *Manufacturing Service Operations Management*, *9*(4), 383-408.

Gittins, J. C. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)*, *41*(2), 148-177.

Grant, R. (1996). Toward a knowledge-based theory of the firm. *Strategic Management Journal*, *17*, 109-122.

Greve, H. (2002). Sticky aspirations: Organizational time perspective and competitiveness. *Organization Science*, *13*(1), 1-17.

Greve, H. R. (2008). A behavioral theory of firm growth: Sequential attention to size and performance goals. *Academy of Management Journal*, *51*(3), 476-494.

Gruber, M., MacMillan, I. C., & Thompson, J. D. (2008). Look before you leap: Market opportunity identification in emerging technology firms. *Management Science*, *54*(9), 1652-1665.

Harrison, J. R., & March, J. G. (1984). Decision making and postdecision surprises. *Administrative Science Quarterly*, *29*(1), 26-42.

Hardwick, J. P., & Stout, Q. F. (1992). Bandit strategies for ethical sequential allocation. In *Proceedings of the 23rd Symposium on the Interface, American Statistical Association: Vol. 23. Computing science and statistics* (pp. 421-4).

Hmieleski, K. M., & Baron, R. A. (2008). When does entrepreneurial self-efficacy enhance versus reduce firm performance? *Strategic Entrepreneurship Journal*, *2*(1), 57-72.

Holland, J. (1975). *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control & artificial intelligence*. Ann Arbor: University of Michigan Press.

Hu, S., Blettner, D., & Bettis, R. A. (2011). Adaptive aspirations: Performance consequences of risk preferences at extremes and alternative reference groups. *Strategic Management Journal*, *32*(13), 1426-1436.

Kanter, R. M. (2006). *Confidence: How winning streaks and losing streaks begin and end*. Crown Business.

Knudsen, T., & Levinthal, D. A. (2007). Two faces of search: Alternative generation and alternative evaluation. *Organization Science*, *18*(1), 39-54.

Lave, C., & March, J. (1975). An introduction to models in the social sciences. *New York, Harper and Row*, 421.

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

Levinthal, D. A., & March, J. (1993). The myopia of learning. *Strategic Management Journal*, *14*(Special Issue), 95-112.

Levinthal, D. A., & Rerup, C. (2006). Crossing an apparent chasm: Bridging mindful and less-mindful perspectives on organizational learning. *Organization Science*, *17*, 502-513.

Makadok, R., & Barney, J. (2001). Strategic factor market intelligence: An application of information economics to strategy formulation and competitor intelligence. *Management Science*, *47*(12), 1621-1638.

March, J. G., & Olsen, J. P. (1975). The uncertainty of the past: Organizational learning under ambiguity. *European Journal of Political Research*, *3*(2), 147-171

March, J., Sproull, L., & Tamuz, M. (1991). Learning from samples of one or fewer. *Organization Science*, *2*(1), 58-71.

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, *2*(1), 71-87.

March, J. (1994). *A primer on decision making : How decisions happen*. The Free Press

March, J. G. (1996). Learning to be risk averse. *Psychological Review*, *103*(2), 309-319.

March, J. (1996). Continuity and change in theories of organizational action. *Administrative Science Quarterly*, *41*(2), 278-287.

March, J. G. (2003). Understanding organizational adaptation. *Society and Economy*, *25*(1), 1-10.

March, J. G. (2010). *The ambiguities of experience*. Ithaca, NY: Cornell University Press.

Nelson, R., & Winter, S. G. (1982). *An evolutionary theory of economic change* Cambridge, MA: Belknap Press.

Nonaka, I., & Von Krogh, G. (2009). Perspective - Tacit knowledge and knowledge conversion: Controversy and advancement in organizational knowledge creation theory. *Organization Science*, *20*(3), 635-652.

Posen, H. E., & Levinthal, D. A. (2012). Chasing a moving target: Exploitation and exploration in dynamic environments. *Management Science,* 58(3), 587-601.

Posen, H. E., Martignoni, D. A., & Levinthal, D. (2012). E Pluribus Unum: Organizational Size and the Efficacy of Learning.

Simon, H. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*, 99-118.

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

Weick, K. E. (1999). Theory construction as disciplined reflexivity: Tradeoffs in the 90s. *Academy of Management Review*, *24*(4), 797-806.

FIGURES

Figure 1: Long-Run Performance Implications of First- and Second-order knowledge

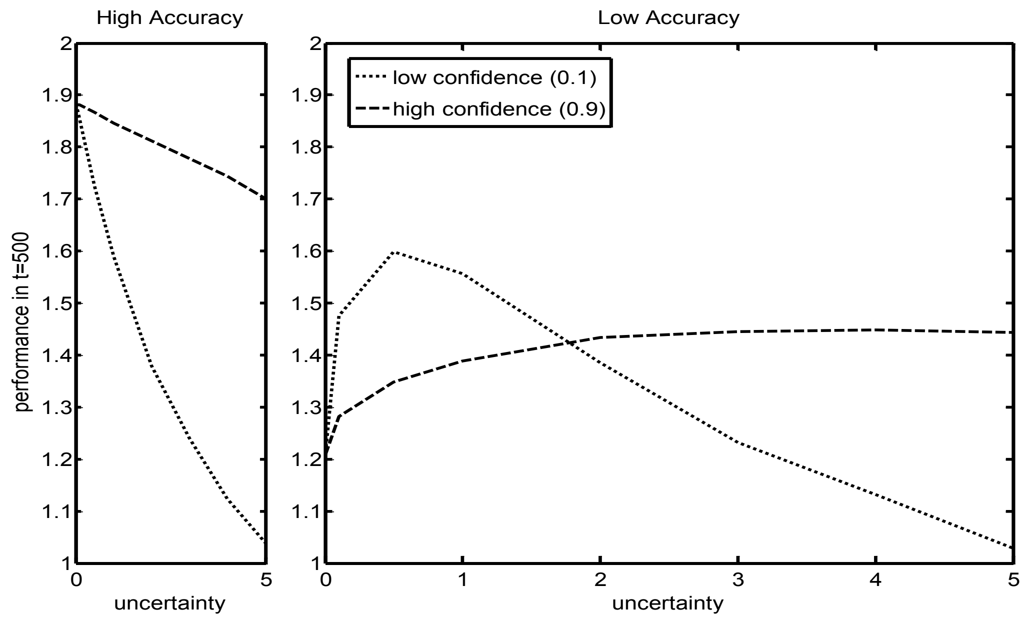

Figure 2: Performance Evolution for Low Accurate Beliefs

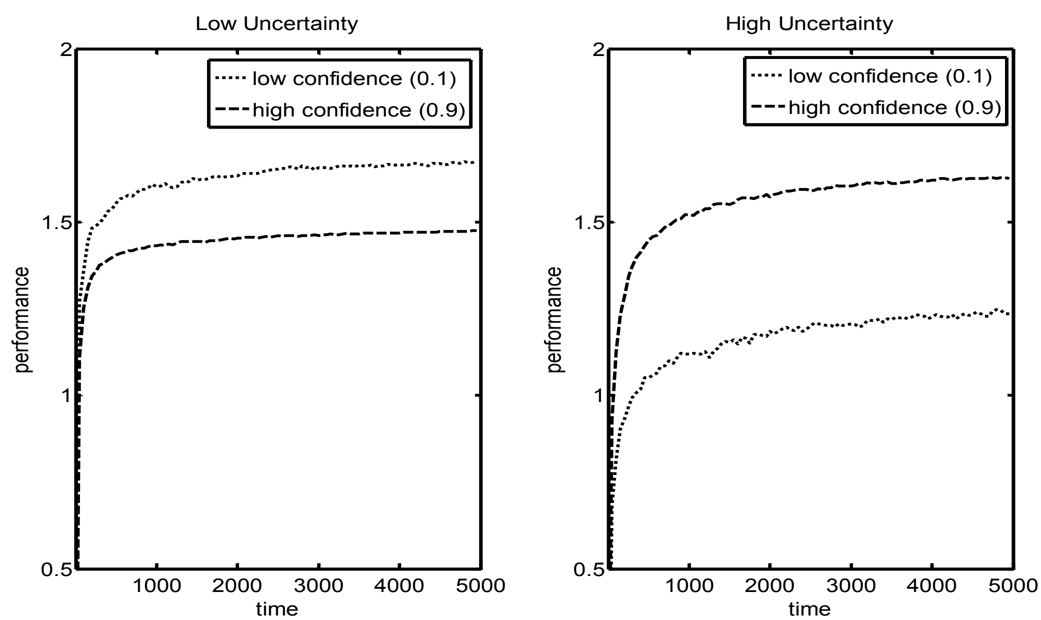

Figure 3: Decomposing the Performance Effect: Exploration and Evaluation Effect

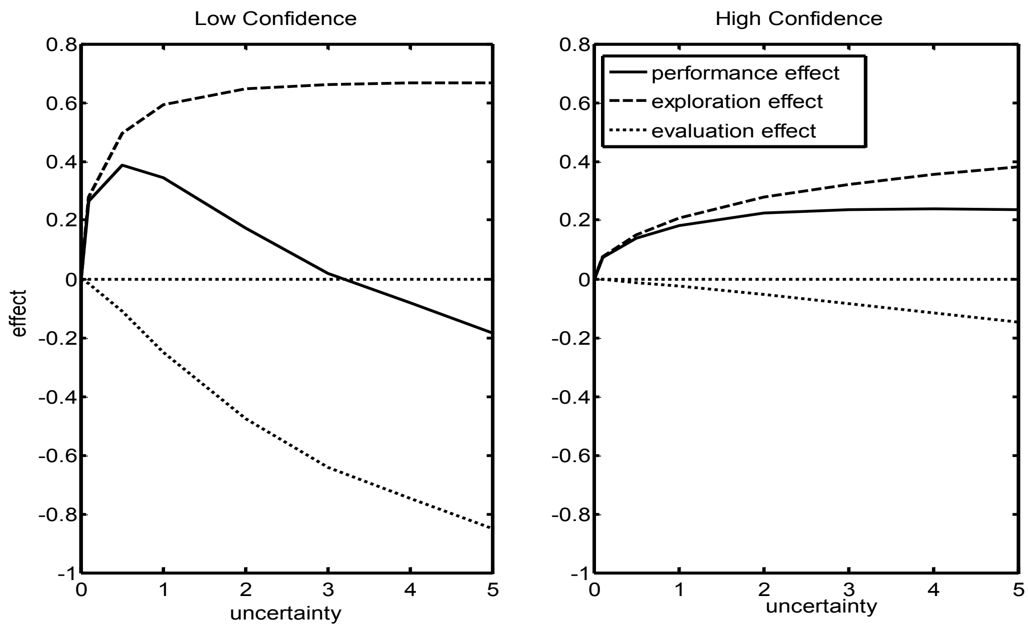

Figure 4: Evaluation Effect

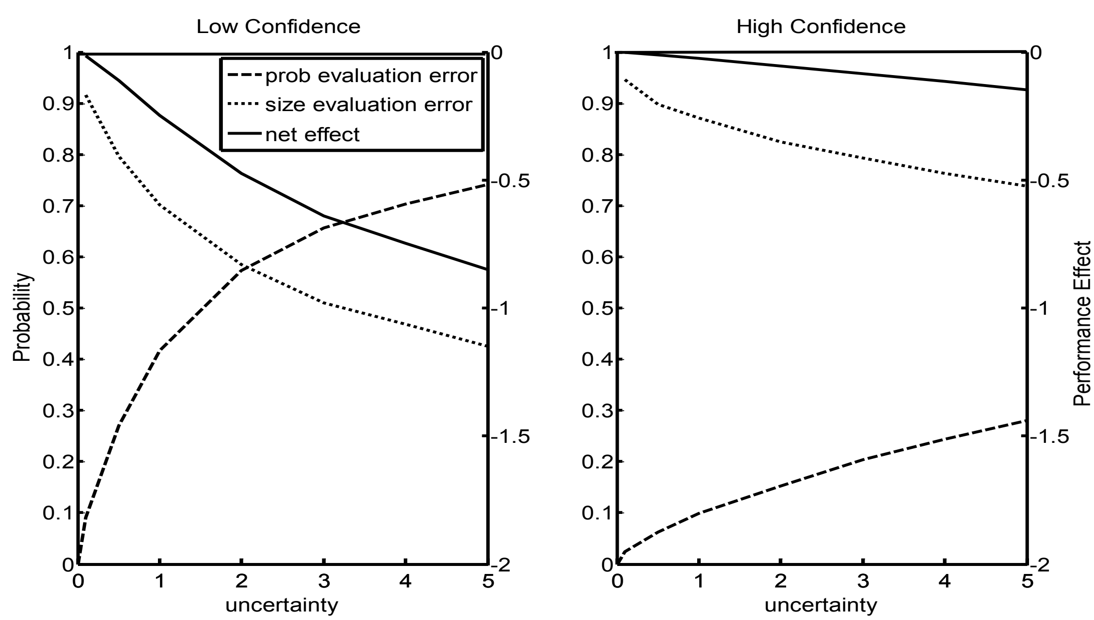

Figure 5: Extent of Exploration

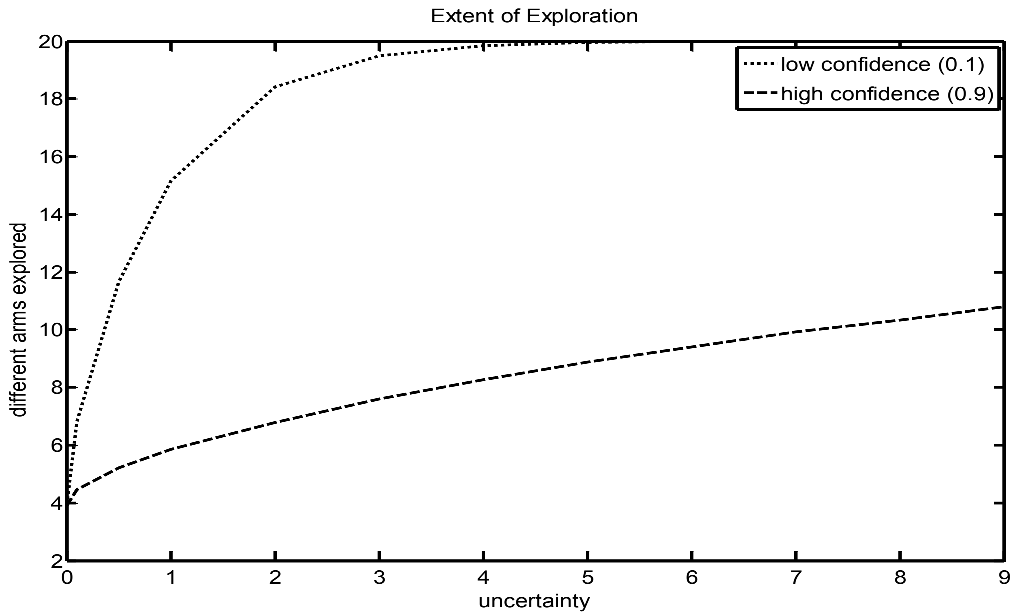

Figure 6: Effect of Exploration