Journal of Neuroscience, Psychology, and Economics
2010, Vol. 3, No. 2, 95–115

1937-321X/10/$12.00 DOI: 10.1037/a0018495

The Neuroscientific Foundations of the Exploration–Exploitation Dilemma

Daniella Laureiro-Martínez
Bocconi University and Universidad de los Andes

Stefano Brusoni and Maurizio Zollo
Bocconi University

What are the origins of the ability to continuously explore novel domains of activity while at the same time exploiting the current knowledge base with increasing efficacy? The conflicting objectives of exploration and exploitation compete for scarce resources, among which managerial attention is possibly the most critical. This paper integrates recent findings on the neuromodulation of attention to provide a foundational step in understanding how the mind of the manager handles the exploration–exploitation dilemma. Also, this paper proposes several possible ways to combine research in neuroscience, psychology, and management to advance our knowledge of the microfoundations of managerial decision-making.

Keywords: microfoundations, exploration, exploitation, attention modulation

Adaptive firm behavior in a diverse and rapidly changing environment requires a trade-off between exploiting known sources of reward and exploring the environment for more valuable or stable opportunities. This trade-off is known as the exploration–exploitation dilemma (March, 1991) and is present at different levels of analysis and different time scales of decision-making. There is no general optimal policy for how to manage the trade-off between exploration and exploitation in nonstationary environments (Cohen, McClure, & Yu, 2007).

From a managerial point of view, the exploration–exploitation dilemma, which is key to many of the challenges faced by organizations, refers to the difficulty faced by organizations and their members in trying to find a balance among competing activities in the context of scarce resources—the need to be efficient to get the most from a current situation, while at the same time exploring possibilities for future improvements. As March (1991) puts it:

> Adaptive systems 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. They exhibit too many undeveloped new ideas and too little distinctive competence. Conversely, systems that engage in exploitation to the exclusion of exploration are likely to find themselves trapped in suboptimal stable equilibria. As a result, maintaining an appropriate balance between exploration and exploitation is a primary factor in system survival and prosperity (March, 1991).

Levinthal and March (1993) similarly argue that "an organization that engages exclusively in exploration will ordinarily suffer from the fact that it never gains the returns of its knowledge" and that "an organization that engages exclusively in exploitation will ordinarily suffer from obsolescence" (p. 105). A narrow search can lead to increasingly rigid cognitive maps and highly specialized competencies that may become core rigidities (Leonard-Barton, 1995). The so called "ambidexterity hypothesis" (Tushman & O'Reilly, 1996) states that the higher the organizational ability to balance ex-

Daniella Laureiro-Martínez, Department of Management, Bocconi University, and School of Management, Universidad de los Andes; Stefano Brusoni, KITeS-CESPRI, Department of Economics, Bocconi University; and Maurizio Zollo, CROMA, Department of Management, Bocconi University.

We thank Nicola Canessa, Stefano Cappa MD, Luis Felipe Orozco-Cabal MD, David Papo and the participants in the seminars at the IRI Summer School, the Bocconi Management Seminars, and the JNPE and SMS Conferences for their comments and suggestions. Any mistakes in this article are the sole responsibility of the authors.

Correspondence concerning this article should be addressed to Daniella Laureiro-Martínez, Via Roentgen, 1, Office 5 B1-12, 20136 Milan, Italy. E-mail: daniella.laureiro@unibocconi.it

ploration and exploitation, the higher the organizational performance.

Since March's (1991) seminal article, the management literature has used the terms "exploration" and "exploitation" in studies of organizational adaptation, organizational learning, competitive advantage, technological innovation, organization design, and organizational survival. However, "an examination of the literature indicates that the answers contained there to the central questions on this subject remain incomplete, at times contradictory, and at best ambiguous" (Gupta, Smith, & Shalley, 2006). The review in Section 2 shows that while the exploration−exploitation literature is extensive and growing, there are gaps that we believe derive from a lack of understanding of the dilemma at the micro level, that is, the individual decision-makers' point of view. There is an intriguing opportunity to better understand this point of view. The recent development of knowledge in the cognitive neurosciences opens, in fact, exciting possibilities to build integrative approaches to understand organizational dilemmas from a micro perspective. That is the overarching objective in this paper: to explore an organizational conundrum from an individual level of analysis, showing how interpersonal variation in the decision-maker's neurological disposition affects the decisional outcomes and, potentially, the performance of a given task connected to that decision.

To do so, the paper is organized as follows: Section 2 briefly presents the main gaps in the organizational literature on exploration–exploitation, focusing on the one that this paper aims to contribute to; Section 3 addresses the microfoundations—that is, the individual origins—of this dilemma, whereas Section 4 contributes by bridging with neuroscience and discussing certain findings in that domain that may help to clarify the roots of the managerial dilemma and suggest ways to cope with it. We then provide an illustration of how the theory we developed might explain the behavior and the outcomes connected to one of the most famous innovators of modern times: Thomas Alva Edison. Section 6 then presents the main challenges involved in the development of an empirical agenda to validate these ideas, offering a full set of suggestions on how to tackle them, and Section 7 concludes with several suggestions for the development of this line of work in future research.

Gaps in the Management Literature on the Exploration–Exploitation Dilemma

A review of the literature on the exploration−exploitation managerial dilemma uncovers several limitations and gaps, which we briefly cover with the aim to focus on and potentially contribute toward resolving one particular gap: variation among individuals in the tendency to respond in an exploitative or explorative way to a given stimulus, and on the ability to shift their responses according to changes in the environmental conditions, have not been explored.

A central gap is that most of the extant research focuses on the structural antecedents to and the effects of involvement in exploration and exploitation on firm performance (Raisch & Birkinshaw, 2008). Few studies delve more deeply into *how* these two activities occur simultaneously. Of course, part of the *how* question has to do with the actual capacities and behaviors of *individual members of the organization*, rather than with organizational arrangements and collective processes. Another key gap is the lack of clarity about the levels of analysis in research on exploration−exploitation. Recent theoretical contributions (Felin & Foss, 2005; Felin, & Hesterly, 2007; Rothaermel & Hess, 2007) identify two main problems with the single-level research approach. First, focusing on only one level of analysis implicitly assumes that most of the heterogeneity is located at that level while other levels are more or less homogeneous. Second, this focus implies that this level is independent of interactions with other lower- or higher-order levels of analysis.

Another main concern with the current literature is the lack of agreement about key elements regarding the definitions of exploration and exploitation. It is not clear from the literature whether exploration and exploitation should be viewed as two ends of a continuum, or as two different and orthogonal aspects of organizational behavior. The central ambiguity in the definitions is whether exploration and exploitation differ in the type of learning or by the presence/absence of learning (Gupta et al., 2006). Table 1 summarizes some of the definitions that appear in key articles on the subject.

Table 1
Definitions of Exploration and Exploitation

Quotes from:	Exploration	Exploitation
March, 1991	Exploration includes elements captured by terms such as search, variation, risk-taking, experimentation, play, flexibility, discovery, and innovation. The essence of exploration is experimentation with new alternatives. Its returns are uncertain, distant, and often negative.	Exploitation includes such things as refinement, choice, production, efficiency, selection, implementation, and execution. The essence of exploitation is the refinement and extension of existing competences, technologies, and paradigms. Its returns are positive, proximate, and predictable.
March, 2006	Pursuit of what might come to be known.	The refinement and implementation of what is known.
Holmqvist, 2004	Exploration is concerned with creating variety in experience, and thrives on experimentation and free association. Variety in experience through search, discovery, novelty, innovation, and experimentation.	Exploitation is about creating reliability in experience, and thrives on productivity and refinement. Creates reliability in experience through refinement, routinization, production, and implementation of knowledge.
Levinthal and Rerup, 2006	"Experimenting with a novel action implies forgoing the use of existing, established practices. In this sense, mindfulness corresponds to exploratory behavior . . ."	". . . and less-mindful behavior is akin to exploitative behavior."
Zollo and Winter, 2002	Exploration activities are primarily carried out through cognitive efforts aimed at generating the necessary range of new intuitions and ideas (variation), as well as selecting the most appropriate ones through evaluation and legitimization processes.	By contrast, exploitation activities rely more on behavioral mechanisms encompassing the replication of the new approaches in diverse contexts and their absorption into the existing sets of routines for the execution of that particular task.
Smith and Tushman, 2005	Exploration is rooted in variance-increasing activities, learning by doing, and trial and error; exploration creates futures that may be quite different than the organization's past.	Exploitation is rooted in variance-decreasing activities and disciplined problem-solving exploitation builds on an organization's past.
Levinthal and March, 1993	"the pursuit of new knowledge of things that might come to be known."	"the use and development of things already known."

In order to avoid confusion, in this paper we adopt a definition very close to the one provided in neuroscience (discussed in Section 4) that admits learning in both exploration and exploitation. We define exploration as the behavior that includes search for alternatives and disengagement from the current task. The simplest form of exploration is random search but other more structured types of search, such as the use of heuristics or explicit algorithms, are also included. As a consequence of this behavior, experimentation, flexibility, discovery, and innovation are shown. We define exploitation as the behavior that helps optimize task performance. When this behavior is present, there is a high engagement with the current task. As a consequence of this behavior, selection, refinement, choice, production, and a concern with efficiency are shown.

The fourth major gap in the received literature on the exploration–exploitation dilemma, and the one we focus on, concerns the role of the characteristics of individual traits as mechanisms underpinning the development of organizational capabilities related to the balanced management of exploration and exploitation activities. The roles of individual and group characteristics were viewed as an important and necessary focus of scholarly attention since the inception of the behavioral theory of the firm (Cyert & March, 1963; March & Simon, 1958; Simon, 1985). March (1991), for example, describes the cognitive limits that individuals encounter when trying to conduct explorative and exploitative processes simultaneously. And Tushman and O'Reilly (1996) argue, the ability to explore–exploit at the organizational level is facilitated by the top-management team's internal processes.

It is only recently, however, that scholars have started to empirically investigate team characteristics that enable organizations to manage both exploration and exploitation (Beckman, 2006; Lubatkin, Simsek, Ling, & Veiga, 2006; McKenzie, Woolf, van Winkelen, & Morgan, 2009; Mom, Van Den Bosch, & Volberda, 2009; Perretti & Negro, 2006; Smith & Tushman, 2005). Gibson and Birkinshaw (2004, p. 223), for example, note the "important role played by senior executives in making an organization context effective and developing ambidexterity" (p. 223). Similarly, Smith and Tushman (2005) explore the integrative mechanisms by which leadership teams might successfully manage the contradictions that arise from structural separation in ambidextrous organizations, and Volberda, Baden-Fuller, and Van Den Bosch (2001) note that "top management explicitly manages the balance of exploration and exploitation by bringing in new competencies to some units while utilizing well-developed competencies in others" (p. 165). And at the group level of analysis, Beckman (2006) finds evidence that the composition of the founding team, and members' prior company affiliations in particular, is an important antecedent to firms' exploitative and explorative behaviors.

Unfortunately, the focus on team characteristics as the antecedents to the development of organizational capabilities in ambidexterity is not matched by studies on the role of managers' individual characteristics or on the ability to make balanced exploration−exploitation decisions. O'Reilly and Tushman (2004), for instance, emphasize the role of ambidextrous managers with "the ability to understand and be sensitive to the needs of very different kinds of businesses" (p. 81). Despite a seeming consensus on the importance of the individual, the state of the art in scholarly work on the individual level of analysis shows a concerning paucity, with the notable exception of Mom, Van Den Bosch, and Volberda (2007). In our view, this is an important concern, as many current "holistic" explanations might capture at least some of what really are effects of variation at the individual level of analysis (Felin & Foss, 2005; Felin & Hesterly, 2007). We explore this further in Section 3.

The Exploration−Exploitation Dilemma in the Manager's Mind

Most studies of exploration−exploitation, starting with March's (1991) seminal work, focus on levels of analysis above the individual. March (1991) analyzes the exploration−exploitation trade-off in the social context of organizations, focusing on two distinctive features: mutual learning in the organization and the individuals involved, and the competition for primacy among organizations. Implicit in his model is the assumption that the balance between exploration and exploitation is based on a turnover process among a mix of individuals with different cognitive characteristics (some more inclined to exploration, others more inclined to exploitation), who achieve a trade-off for the whole organization.

The mechanism based on the turnover of "cognitively specialized" (and inflexible) managers to achieve a balance between those predisposed to exploitative behavior and those predisposed to explorative behavior, however, is clearly not the only one at the disposal of the organization (the alternative, of course, is to gather a group of managers who can think and act in both modes with relative ease), and probably not even the optimal one. First of all, having a group of cognitively specialized decision-makers will simply turn the decision into a political battle between the two factions, with results driven by the relative size and political weight of the two factions, rather than to rational choice. Second, even when the decision can be efficiently allocated to "exploration-minded" or to "exploitation-minded" managers, this decision itself requires a decision-maker (e.g., the CEO) with a significant amount of cognitive flexibility to recognize the advantages and disadvantages of the two alternative allocations of the decision-making responsibility. Essentially, the problem (and the solution, via cognitive flexibility) would still be there, but would be upgraded to the cognitive profile of the person at the top of the organization, who assigns the problem to the cognitively specialized groups.

The alternative scenario of having cognitively nonspecialized (and flexible) managers might thus be superior to the one described above, since it would reduce the likelihood of political fights and suboptimal decisions and would not require the presence of a "higher order" decision-maker to assign the responsibility of the solution to "cognitively specialized" groups of managers. As O'Reilly and Tushman (2004) recognized, "one of the most important lessons is that ambidextrous organizations need ambidextrous senior teams and managers" (p. 81).

We thus propose that to achieve a better understanding of how both exploration and exploitation can be conducted, we need an in-depth examination of the microfoundations of organizational learning. The micro—that is, individual—level of analysis may in fact account for an important amount of heterogeneity in decisional and performance outcomes, and should be explicitly studied:

to fully explicate organizational anything—whether identity, learning, knowledge or capabilities—one must fundamentally begin with and understand the individuals that compose the whole, specifically their underlying nature, choices, abilities, propensities, heterogeneity, purposes, expectations and motivations. While using the term "organizational" may serve as helpful shorthand for discussion purposes and for reduced-form empirical analysis, truly explaining the organization (e.g., existence, decline, capability or performance), or any collective for that matter, requires starting with the individual as the central actor (Felin & Foss, 2005).

Mom et al. (2007) were, to the best of our knowledge, the first to analyze the exploration–exploitation dilemma at the individual level of analysis. They explored the influence of knowledge flows on the manager's explorative or exploitative activities, recognizing that one of the most promising avenues for future research is "measuring exploration and exploitation at the managerial level of analysis using objective measures" (p. 927).

Along these lines of discourse, Section 4 discusses recent findings on the neuromodulation of attention, which, we argue, is at the core of the human ability to shift from one learning mode to another. We provide some suggestions on how to build on the concepts of situational uncertainty, utility perception and attention focus to further investigation of organizational exploration and exploitation.

A Neuroscientific Approach to the Exploration−Exploitation Dilemma

In line with our aim to examine the microfoundations of the exploration−exploitation dilemma, we searched for work on the cognitive processes underlying the exploration−exploitation trade-off. Recent work on the neuromodulation of attention proposes that interactions between the orbitofrontal cortex (OFC), the anterior cingulate cortex (ACC), and the locus coeruleus (LC) (see Figure 1) may modulate attention and thus balance exploration−exploitation (Aston-Jones & Cohen, 2005; Cohen, Aston-Jones, & Gilzenrat, 2004; Usher, Cohen, Servan-Schreiber, Rajkowski, & Aston-Jones, 1999).

Traditionally, the locus coeruleus–norepinephrine (LC-NE) system was considered to be implicated solely in arousal. However, multiple recent findings suggest that this system plays a more complex and specific role in the control of behavior by contributing to the optimization of behavioral performance (Sara, 2009).

Aston-Jones, Rajkowshi, Kubiak, and Alexinsky (1994) observe that the LC shifts between two operating modes: the phasic and the tonic. In the former, LC cells exhibit phasic activation in response to the processing of task-relevant stimuli, but display only a moderate level of tonic discharge. This mode is consistently associated with enhanced attention focus, generating "exploitative" behavior, defined as behavior that optimizes and achieves high levels of task performance (Aston-Jones & Cohen, 2005). It is important to note here that the analysis is at the individual level. Exploitative behavior translates into a high level of engagement with the current task. As a consequence, behaviors that show refinement in the selected actions and a concern for efficiency are demonstrated.

In the tonic mode, LC cells do not respond phasically to task events, but exhibit higher

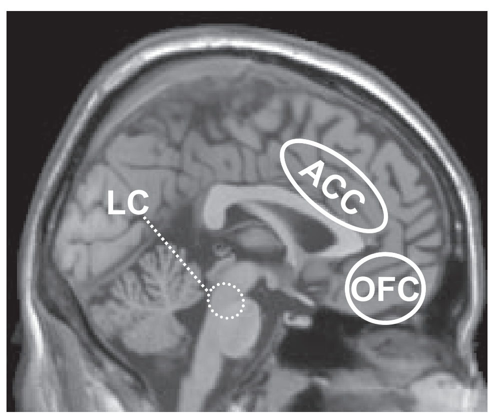

Figure 1. LC, OFC and ACC location in the brain.

levels of ongoing tonic activity. Exploration is defined as the behavior shown when there is search for alternatives and disengagement with the current task (Aston-Jones & Cohen, 2005). This mode is associated with explorative behavior because it corresponds with poor performance on tasks that require focused attention, and with increased distractibility. The simplest form of exploration is random search, but exploration also includes more structured types of search, such as the use of heuristics or explicit algorithms. This behavior demonstrates abilities for experimentation, flexibility, discovery, and innovation.

It should be noted that the definitions of exploration and exploitation provided by neuroscientists (Aston-Jones & Cohen, 2005) are compatible with those in the management literature (see Table 1) and with March's (1991) definition: "The essence of exploitation is the refinement and extension of existing competences, technologies, and paradigms. Its returns are positive, proximate, and predictable. The essence of exploration is experimentation with new alternatives. Its returns are uncertain, distant, and often negative." (p. 81).

The findings on the functioning of the LC and its consequent type of behavior are based on two pieces of evidence. The first comes from experiments on monkeys, which transitioned from a phasic to a tonic mode and then reversed when the new target was identified. This transition requires that LC has the relevant information to determine when to switch between phasic and tonic modes, an important aspect that we address in the section on perception of utility (Usher et al., 1999). The second piece of evidence derives from studies of humans measuring pupil diameters—a good proxy for LC activity—and functional MRI (fMRI) experiments (see Beversdorf, White, Chever, Hughes, & Bornstein, 2002; Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006; and Sterpenich et al., 2006; among others). Usher and colleagues (1999) also develop a biophysically plausible model of LC functioning that accounts for transitions between the phasic and tonic modes in terms of a single physiological variable (coupling between LC cells) and explains the impact of these shifts on task performance. In brief, the model suggests that the phasic mode favors exploitation by releasing norepinephrine (NE) when a task-relevant event occurs, thereby facilitating the processing of that event, while sustained release of NE in the tonic mode indiscriminately facilitates the processing of all events irrespective of their relevance to the current task, thereby favoring exploration.

Viewed from the perspective of attention, the LC phasic mode supports the current control state (exploitation), while the LC tonic mode provokes a withdrawal of control from the current task, favoring the sampling of other behavioral goals (exploration). These changes between phasic and tonic modes are the basis for an understanding of the exploration–exploitation dilemma from an attention perspective. When the utility derived from a given behavior is low in comparison to expectations, flexibility to change the attention focus, and thus the behavior from exploitation to exploration, is needed to explore the environment and sample different behaviors until new sources of reward are discovered. This is the role played by different modes of activity in the LC-ACC/OFC system. Aston-Jones and Cohen (2005) observed that the LC shifts between two distinct operating modes, and that these shifts change attention and then behavior. The LC phasic mode supports the focus of attention on the current control state (exploitation), while the LC tonic mode provokes a withdrawal of control from the current task, thus favoring broader attention and the sampling of other behaviors (exploration).[1] The phasic LC responses facilitate context-congruent behavioral responses (exploitation) and the tonic mode of LC facilitates sensitivity to different stimuli and the execution of a broader class of behavioral responses (exploration).

These neuroscientific findings on changes in attention scope contribute to our understanding of the management exploration–exploitation dilemma from an attention perspective. The idea that broad attention is important in situations that are dynamic, ill-structured, ambiguous, and unpredictable is acknowledged in the management literature (Levinthal & Rerup, 2006; Weick & Sutcliffe, 2006). In the opposite scenarios, the economies of narrow attention are more appropriate and, through their reliance on routines, can be cost-efficient (Nelson & Win-

[1] We should note here that, whereas we have described the phasic and tonic modes as distinct, they likely represent the extremes of a continuum of function. For expository purposes it is more useful to distinguish them at the two extremes.

ter, 1982). As a consequence, the higher the uncertainty in a situation, the higher the likelihood that broad attention—and thus explorative behaviors—will lead to better performance.

However, these explanations still leave out the important question of what information the neural system uses to determine whether it should exploit (LC phasic mode) or explore (LC tonic mode). Studies have found that the brain computes the perceived utility in a particular situation, compares that with its expectation levels, and drives shifts between LC phasic and tonic modes by influencing simple physiological parameters (Usher et al., 1999). When low expected utility is perceived, the broad mode of attention (tonic mode of the LC functioning) is activated. When high utility is perceived, the focused mode (the phasic LC mode) is activated. Studies of humans (and neuronal records in primates) show that the frontal cortex plays an important role in the evaluation of utility. Different areas of the brain are found to be related to the assessment of rewards and costs. A large number of neuroimaging studies, involving diverse experiments, have examined brain responses to reward stimuli. They consistently identify a common set of neural structures that are activated in response to these stimuli, mainly the OFC, the ventral striatum, and the ventromedial prefrontal cortex. The OFC has been implicated in hedonic experience across all sensory modalities (Rolls, 2000). Of specific interest are areas in the striatum and the OFC that are particularly responsive to rewards and which change, accumulate, or are learned over time (Knutson, Fong, Bennett, Adams, & Hommer, 2003; Koepp et al., 1998; Murray, O'Doherty, & Schoenbaum, 2007).

Attempts have also been made to identify the brain areas activated by cost-related issues. The anterior cingulate cortex (ACC) is traditionally considered to be directly responsive to aversive interoceptive and somatosensory stimuli, and particularly to pain (e.g., (Peyron, Laurent, & García-Larrea, 2000). Recent neurophysiological studies on monkeys and on humans have consistently demonstrated that ACC is strongly responsive to negatively valenced information, such as performance errors, negative feedback, monetary losses, and even social exclusion (Eisenberger, Lieberman, & Williams, 2003), and also to task difficulty and decision-making conflicts (Botvinick, Cohen, & Carter, 2004). Therefore, there is much evidence to suggest that the OFC is involved in reward evaluation, whereas the ACC is responsive to a variety of negatively valenced signals (Aston-Jones & Cohen, 2005; McClure, Laibson, Loewenstein, & Cohen, 2004; Montague, King-Casas, & Cohen, 2006; Tom, Fox, Trepel, & Poldrack, 2007). These results point to the existence of a strong relationship between situation uncertainty and utility perception. Thus, risk-averse individuals, all else being equal, will find higher utility in less uncertain situations and, as a consequence, will act under the phasic LC mode, showing more exploitative behavior. In highly uncertain situations, the opposite will occur. Individuals will find less utility in such situations and will shift to a tonic LC mode, acting in a more explorative way.

That perceiving a high utility reduces the attention to search and exploration, and that the perception of low utility promotes search and explorative behavior, are fundamental tenets of the behavioral school (Cyert & March, 1963; March & Simon, 1958), a departure from standard neoclassical thinking, which assumes constant investment in search and exploration, independent of the current task. The "satisficing" behavior assumption has been widely tested in and supported by management studies over the decades (see Grève, 2003, for a recent review), and by evidence on the impact of past performance on investment in attention and learning by firms (Bateman & Zeithmal, 1989; Simon, Houghton, & Aquino, 2000; Thomas, Clark, & Gioia, 1993).

In addition, the perception of uncertainty has been proposed as an important factor affecting the evaluation of a given situation (McClure, Gilzenrat, & Cohen, 2006; Yu & Dayan, 2005). It seems that our brains distinguish between expected and unexpected uncertainty. Two important neuromodulators—acetylcholine (ACh) and norepinephrine (NE), respectively—signal expected and unexpected sources of uncertainty. When exploiting, if prediction errors are higher than expected, the current strategy should be revised and we should explore. If, on the contrary, the prediction errors can be accounted for in terms of expected uncertainty, the exploiting strategy should be maintained. In general, taking into account that individuals are risk-averse, then the higher the unexpected uncertainty perceived in a situation or a problem, the more difficult it will be for the individual to understand the outcome of the situation/problem and so the less the utility from the situation.

The neuroscientific findings summarized above provide the basis for an understanding of what underlies the exploration–exploitation dilemma. These findings also help to bridge some of the gaps described in Sections 2 and 3. The neuroscientific definitions of exploration and exploitation, and the discovery of the neurological mechanisms underlying the shift between the two attentional states, help to resolve the debate over whether exploration and exploitation are to be viewed as positions on a continuum or as orthogonal situations. At the individual level of analysis, they are clearly orthogonal, since at any given moment an individual cannot be in both the phasic and tonic modes of LC functioning. At the organizational level, however, members of a group of individuals might be functioning in different modes, which means that the group overall will be working on a continuum between a completely phasic mode (all focused on the current task, i.e., purely exploitative mode) and a completely tonic mode (all focused on exploration). Of course, any group typically works at a position on the continuum located somewhere between the two extremes, but this position is important to the prediction of collective behavior (see Figure 2). It is important that the (discrete) shifts in LC operating modes of individuals in the group over time will cause the group position on the continuum to constantly shift.

In the next section, we offer a managerial illustration of how these neurological mechanisms might influence the activities of groups involved in research and development (R&D) work.

Exploration–Exploitation and the "Wizard of Menlo Park"

How are the ideas presented in this paper reflected in a real case? In this section, we argue that an organizational process of strategic relevance, such as a product innovation process, can be decomposed in subactivities (explorative–exploitative behaviors) that are largely influenced by a decision-maker's attention focus (broad and narrow), which we expect to in turn originate from neuromodulatory mechanisms guided by the LC.

We will exemplify the ideas we propose using the well-documented case of Thomas Alva Edison—one of the most famous inventors of all times—and the main events in one year of work on the microphone at Menlo Park, one of the first facilities entirely dedicated to R&D

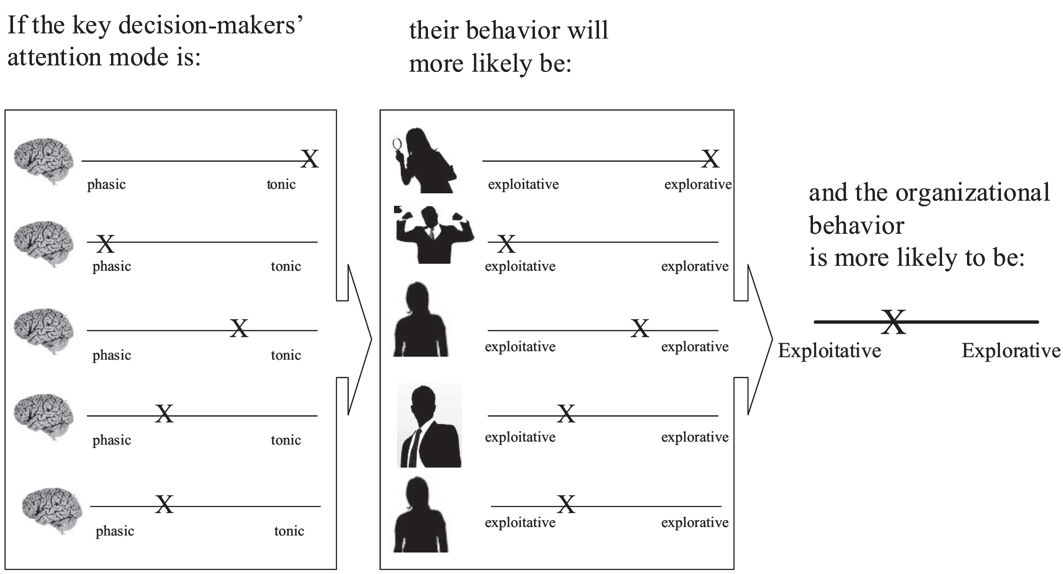

Figure 2. Key decision-makers' attention and organizational behavior.

activities. Edison serves as a good case for illustrating the ideas presented in this paper relating to micro- and macrolevels of analysis since he was not the "lone genius" but instead a "collective noun and means the work of many men" (Lindgren, 1979, p. 17 cited in Swedberg, 1993). While Edison was clearly at the head of the Menlo Park operations, they were the result of the collaboration and work of many individuals: the Park "apparatus" served to promote a huge number of inventions ("a minor invention every 10 days and a big thing every six months or so" [Lindgren, 1979, p. 17 cited in Swedberg, 1993]), generating more than 400 patents in its 6 years of operation. There is agreement on the fact that Edison was a relentless innovator or, closer to our argument, a "meta inventor." He made Menlo Park the cornerstone of modern industrial research. The Park was the first industrial laboratory concerned with both creating knowledge and controlling new knowledge application. It is interesting to apply the ideas we have integrated so far to the context of an R&D lab, where not only research was done and many new inventions were discovered, but also the development, polishing, protection, and selling of the inventions was done.

Menlo Park demonstrated the ability to successfully combine both explorative and exploitative activities. Different inventors inside Menlo Park were involved in both explorative and exploitative activities, with Edison leading their work in *both* cases. Edison's high involvement in the different activities in Menlo Park may have had some negative consequences for the development of the incubator. On the other hand, his role exemplifies how the way that leaders attend to the particular problems of the innovation process, facing the different decisions required to handle both explorative and exploitative challenges, contributes to the success of the overall innovation process.

Not only was Edison working on many different projects at once, requiring different modes of operation in his attention system, but he also showed the ability to continuously shift between the two cognitive modes within the context of a single process. On the one hand he could be very focused on advancing knowledge development toward the solution of a very specific problem. On the other hand, he also recombined specialized knowledge from different fields to generate new, broader subfields and more general knowledge.

How did he switch from one mode to the other? In different innovation cases there is evidence of the importance of specialization as domain-specific expertise (Kaufman & Baer, 2006). Specific expertise is usually the foundation for building innovation. In addition, however, in Edison's case (as in other cases studied by Kaufman and Baer), general expertise played a key role, along with the ability to recombine knowledge from different fields, or to apply it from one field to another. As stated by Hargadon and Sutton (2000), Edison and the people in his lab had the ability to "move easily in and out of separate pools of knowledge, to keep learning new ideas, and to use ideas in novel situations" (p. 161). Like Alexander Graham Bell, Samuel Morse, Henry Ford, and others, Edison did not advance science in the way specialists do. These scientists instead focused on and developed an in-depth knowledge about the specific issues (problems, components, etc.) of their inventions, while also broadening and bringing different streams of scientific discoveries into practical devices and systems (Skrabec, 2006). These serial innovators were sometimes able to *explore*, broaden their attention, be creative, recombine and bring together different ideas and also *exploit*, narrow down their attention, concentrate, and focus on solving a specific problem.

To illustrate, a detailed study by Carlson and colleagues (Carlson, 2003, p. 155–156) shows how Edison's work on the microphone (the carbon button transmitter for the telephone) can be summarized as a series of contrasting behaviors in which "During certain periods, Edison varied his lines of research, and then at particular moments, he appears to have selected one line for further development" (p.156). His inventive patterns can be seen as characterized by "explorative" periods (when he played out different lines of investigation) following by "exploitative" periods when he selected and focused by either choosing one of the lines or by recombining the most promising lines.

Figure 3 summarizes the main activities Edison undertook while developing the microphone for the telephone during 1877. The ovals show the moments when he chose which line to proceed with (Diagram adapted from Carlson, 2003, p. 156). For example, in late April 1877,

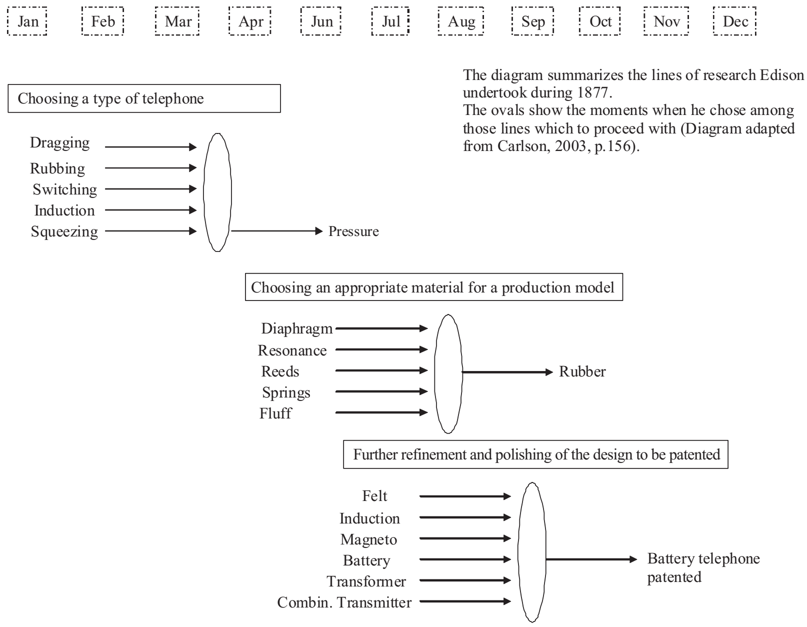

Figure 3. Edison's lines of research for developing the telephone.

Edison was experimenting with different lines (i.e., dragging, rubbing, etc.). He then chose to focus on one option (i.e., the pressure telephone) to further improve it. Again, in September of that year, after having experimented with different production models (i.e., resonance, reed, etc.), he decided to focus on one option. In this case he did not actually chose one, instead combining the most promising results from different lines and focusing on developing a rubber tube production version.

As documented in Carlson (2003, p. 152) this particular situation can be analyzed as a specific case in point:

"Edison or his associate James Adams substituted points of plumbago (i.e., graphite) for the disks on his 'squeeze' telephone. These telephones seemed to work better than previous versions, leading Edison to think more carefully about using points. In particular, he now considered using four high-resistance points pressing on the diaphragm with varying degrees of force. Edison noted an inverse relationship between the mechanical force and the electrical resistance that the resistance increased as the force decreased- and drew on this observation to construct a pressure telephone in April 1877."

A schema for interpreting this example in the light of the ideas proposed in this paper is presented in the Figure 4. The concentric ovals signal the different sublevels at which the exploration−exploitation dilemma can be seen in the telephone example. At the center we present the more micro level we have included in our ideas (the LC mode), and the most external oval represents the more macro level we have presented (the behavior at the organizational level). In the telephone example, Edison started with an exploratory behavior searching for alternatives (in this case five different ones) and experimenting (adding the graphite points). Edison got positive feedback from the environment, obtaining a good outcome (better working telephones), which increased the utility and lowered the uncertainty he perceived in the problem at hand. Perceiving a higher utility, the focused mode of LC functioning (phasic) arose (see first central box to the left). Edison started focusing his attention ("think more carefully about using points. . ." p.156) on one alternative to improve it. He narrowed down from five alternatives to the apparently best one to further exploit the idea, develop and improve it ("he

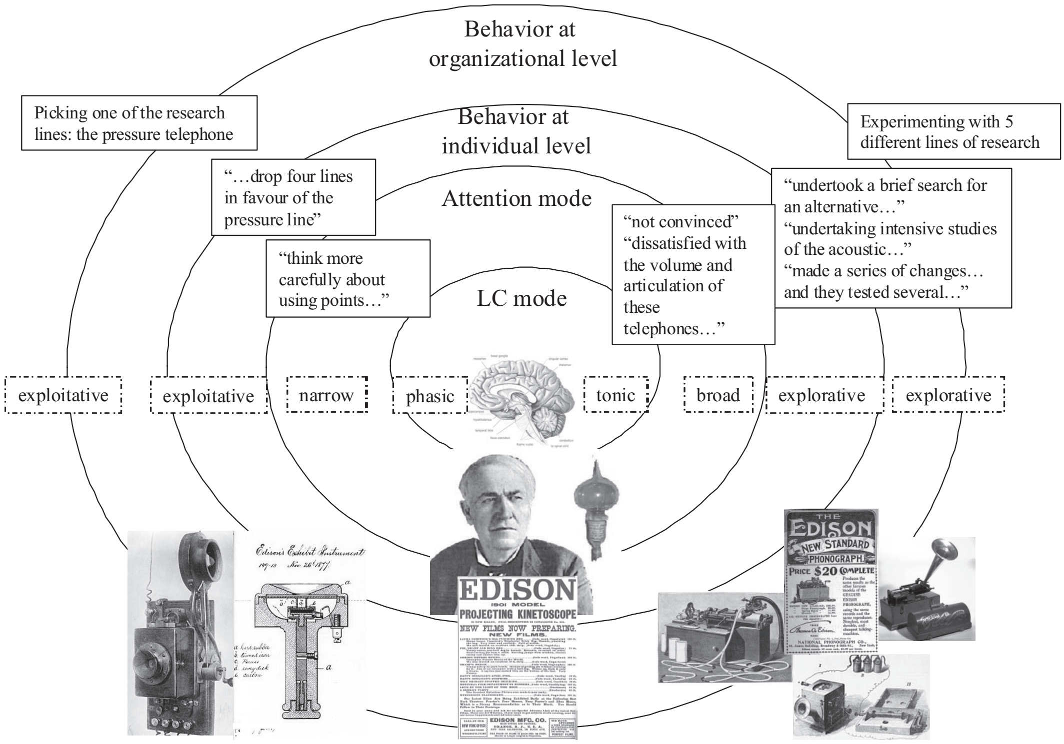

Figure 4. Example: Edison and the telephone invention.

chose to drop four lines in favor of the pressure line" p. 156).

This shifting between exploring and exploiting was repeated several times in the telephone invention and in several others. As can be seen in Figure 4, as a result of the "squeezing" alternative, Edison focused his research on the pressure telephone. His shifts between exploring and exploiting continued during the summer of 1877 (in Figure 4 see the box beginning with "not convinced"). After deciding to focus on the pressure telephone, Edison and his associates faced a different problem: the quality of the acoustic components was not satisfactory for them (low utility perceived). Consequently, he decided to undertake a search for better components (perceiving a low utility, he was under the tonic LC mode and so broadened his attention). The lab then extended their investigation to different acoustic components (material search, diaphragm, resonance cavity, reeds, springs, fluff) that were studied and tested (explorative behaviors). Once higher standards were reached (and so a high utility perceived) again a narrowing-down process appeared in which Edison combined the most promising results from the different lines of research to create the rubber tube production version of the telephone.

If we think about applying the schema in Figure 2 to interpret—somewhat liberally—Edison's case, we can see that if we take a static picture at a moment when Edison and his associates were exploring different lines, we may see how their broad attention (resulting from a tonic LC mode) reflected an explorative behavior at the individual level and aggregated also at the organizational level. At other moments (as, e.g., when narrowing down to one option among the many explored) the more focused attention (resulting from a phasic LC mode) showed exploitative behaviors.

It is important, for our argument, there is no evidence that Edison relied on different people to allocate exploration- or exploitation-oriented tasks, but that all the people involved in his labs tackled both types of challenges. Having individuals capable of shifting easily from exploitation to exploration and back also allows a group to do so: aggregated at the organizational level, the combination of explorative and exploitative behaviors is more likely to lead to an adaptive behavior that swiftly moves between the exploring−exploiting modes.

The Challenge of Measurement and Empirical Validation

In applying neuroscientific findings and techniques to management problems, we need to be aware of some key epistemological differences between these research programs. Neuroscientists, and neuroeconomists for the large part, are interested in the average effects of their experimental manipulations in identifying the neural correlates of a given process (e.g., a utility evaluation) triggered by a homogeneous stimulus (e.g., increased levels of uncertainty) and producing a set of possible, observable, decisional outcomes (e.g., a type of decision in a game-theoretic scenario). Consequently, the pursuit of these types of research questions and the experimental designs generated tend to neglect individual variations around the mean response to the stimulus, and in fact consider them a nuisance (Frederick, 2005), since the larger the variance the weaker are the mean-based results.

However, as Lubinsk1 and Humphreys (1997) note, a neglected aspect does not disappear because it is neglected, and there is no good reason for ignoring the individual variation around a mean response, especially if we have good logic to expect an important causal link between managerial decision-making and organizational performance, or, as in this paper, between the allocation of attention as an antecedent of decision-making behavior and its consequences in organizational behavior and performance. As neuroscientists do, we should care about what causes the average to work in a certain way, but our focus must be on the explanation of the differences: what might lead certain individuals and their organizations to display diverse reactions to similar stimuli, and consequently different levels of performance? This basic difference in epistemological approaches implies some challenges to management researchers who need to adjust some of the traditionally adopted neuroscience research methods in order to comply with their objectives.

A first challenge is that testing for differences requires a larger sample and implies higher costs in terms of time, laboratory access, and data processing. Also, the type of participants is different: many applications of neuroscience to economics, marketing, and finance use data from college students. However, if we want to understand differences in managerial decision-making behavior, it is necessary to sample managers, executives, and entrepreneurs of different kinds, with the aim of tracing differences in the neural correlates that antecede their behavior, and also to eventually be able to correlate the results on individual characteristics with those of the organizations or groups over which they exert significant influence. Differences in age and experience may shape the brain structure and the activations that can be observed. One interesting setting for testing our ideas would be among the entrepreneurs of small family enterprises. In those organizations, as described by Bodmer and Vaughan (2009), the limited resources will possibly increase the impact on the organizational moves of what the entrepreneur attends to and the cognitive maps he does develop. Finally, replication is especially important in managerial studies since contextual effects related to task features, firm and sector characteristics, cultural traits and institutional conditions may alter the way individuals make decisions following the same set of stimuli, and the way organizations consequently perform.

Another distinctive feature of management scholarship requires researchers to go beyond uncovering what the effects of certain abilities might be, and to attempt to understand how these cognitive abilities can be developed and used successfully to improve organizational performance, so that abilities—say the flexibility to shift attentional mode as soon as unexpected uncertainty levels change—can be diffused to benefit organizational outcomes. One of the key principles of behavioral neuroscience, in fact, is that experience can modify brain structure long after brain development is complete. Brain plasticity refers precisely to the brain's ability to change structure and function. Experience is a major stimulant of brain plasticity and works by producing multiple, dissociable changes in the brain including increases in dendritic length, increases (or decreases) in spine density, synapse formation, increased glial activity and altered metabolic activity (Kolb & Whishaw, 1998). Research on humans produced the interesting result that "the plasticity of the nervous system remains throughout the life span and extends well into old age" (Taub, 2004). If the antecedents to certain managerial abilities are identified, it may imply that it might be possible to modify brain structure through different types of exercises and training, and thus gain abilities relevant to improving managerial functioning following different experiences and at different ages.

To design a study for testing the neuroscientific findings in a managerial setting it is necessary to create or adapt experimental tasks that cover a series of steps corresponding to the different constructs illustrated by the neuroscientific findings. An ideal option would be to obtain direct measures of the managers' attention focus while facing different real-life decisions and to correlate such measures with the performance obtained out of those decisions (both at a purely individual level and at an organizational level) both in the short- and the long-term. However, given the difficulty (or impossibility) of measuring managers in the real context with the actual brain imaging tools, a good viable proxy would be to correlate the performance obtained, for example, in managerial decision-making simulations facing the exploration–exploitation dilemma with the precise attention focus measures obtained while completing different neuropsychological tasks in a lab context. There are different alternatives that researchers could follow to measure the ability of making decisions regarding exploring–exploiting, some of which are also compatible with brain imaging techniques that could allow researchers to measure not only performance in the exploration−exploitation decisions (the observed behavior) but also the neural correlates of such decisions.

We now turn to discuss how each of the key constructs in this paper could be actually observed with the context of one specific task, which can be administered using fMRI techniques: the gambling task (Daw et al., 2006).

To understand the antecedents of the decision-making ability related to managing the exploration−exploitation dilemma, we have proposed four constructs based on the findings in neuroscience. The first construct is the level of uncertainty connected to a given decisional situation. The first step in the experimental design is thus to evaluate the manager's perception of the uncertainty of the outcomes in the task. This perception of uncertainty translates into a utility assessment, the second construct. Different parts of the brain intervene in utility assessment and, as we have shown, the ACC and OFC play important roles. Depending on the manager's assessment of utility in the current task, an attention mode arises, the third construct. As already explained, the LC plays a fundamental role in modulating the attention mode according to the perceived utility. As a consequence of the attention mode, the manager will then show a behavior, the fourth construct. In the case of broad attention (LC active in tonic mode), explorative behaviors arise and managers will act (or will propose solutions) in ways characterized by experimentation, flexibility, discovery, and innovation. In the case of focused attention (LC active in phasic mode), exploitative behaviors will occur and managers will act selectively, according to refinement of current processes and efficiency in the current task. If the behavior matches what the situation demand (e.g., high uncertainty matched with explorative behavior and low uncertainty with exploitative behavior), higher performance can be expected.

It is important that the ability to balance exploration−exploitation through flexible management of the situational requirements to achieve the appropriate attentional response can only be assessed if the decision process is replicated under stable contextual conditions. This adaptive process at the individual level can be linked, and the link empirically tested, to the organizational ability to balance the exploration–exploitation dilemma.

The four constructs we propose can be measured in an experimental context. These experiments involve two types of data. First, they require data derived from behavioral measurements during experimental tasks (e.g., response times, performance, etc.) and from questionnaires and interviews. Second, they require data on brain functioning, which can be obtained during experiments by using brain imaging techniques.[2]

The particular technique to be used can be selected depending on the properties of brain–behavior association to be observed. If the study requires a high spatial resolution, fMRI and PET will be required. If the study requires high temporal resolution, EEG or MEG would be suitable. A study that requires both high temporal and high spatial resolution could use a combination of these techniques (such as fMRI or PET with EEG/MEG).

Gambling in the MRI Machine

An experimental task that could assess the exploration–exploitation dilemma at the individual level and is compatible with the brain imaging techniques just exposed is the gambling task as adapted by Daw et al. (2006). In this experiment, participants play by choosing among four slot machines, to win as many points as possible. They are faced with the classical exploration–exploitation dilemma in a changing environment context. During the experiment different characteristics on the machines are manipulated (payoff average, uncertainty of returns, etc.) and participants must choose whether to continue to play on a particular machine or explore new possibilities in the hope of earning more points. While the individual is playing, her or his brain can be scanned (e.g., using fMRI) to obtain measures for each of the four constructs developed above. Of course, the task could be done in a normal PC and obtain only the behavioral performance measures of exploration–exploitation decision-making, depending on the strategies and payoffs obtained by the participant. Also important, the software controlling the game allows a good level of manipulation of the uncertainty level. In addition, the utility perception and the attention mode can be measured using brain imaging techniques. Most management studies proxy for attention by time allocated to an activity (e.g., the famous study by Mintzberg, 1973, whose results were confirmed by Kurke and Aldrich, 1983). Other studies proxy for attention by the number of times a person refers to a certain issue (e.g., the number of sentences in a shareholder letter referring to a specific element (D'Aveni and MacMillan, 1990; Barr, Stimpert, & Huff, 1992). However, the ideas we propose do not require a measure for general attention; rather, we want to differentiate between two types of attention—broad and focused. The attention mode can be assessed by observing LC functioning using a brain imaging technique (e.g., the participants can play the task while lying in the MRI scan) or by measuring pupil diameter, found to correlate remarkably well with LC tonic activity (Gilzenrat, Cohen, Rajkowski, & Aston-Jones, 2003). The diameter of an individual's pupils changes under various conditions. For instance, the pupil diameter of someone who is thinking increases; the pupil of a tired person shrinks. These effects have been proven through a number of psychology (Shinoda & Kato, 2006) and neuroscience experiments (Gilzenrat et al., 2003).

Finally, behavior can be observed based on individual choices. In this way, one task could be used to measure the decision-making performance when facing the exploration–exploitation tradeoff, and for manipulating and measuring the different constructs that we have illustrated and that affect such performance (uncertainty, perceived utility, attention focus). Table 2 summarizes how each construct can be defined and measured in an experiment such as the gambling task.

This gambling task could be complemented with other alternatives, such as simulation, decision-making vignettes and self-reported scales, to have a more reliable measure. A treatment of all the available alternatives using these techniques is out of the scope of this article, but the authors will be pleased to provide a synopsis and an assessment of results from ongoing empirical work to all interested scholars.

[2] The basic set of techniques used to generate neurological images is electroencephalography (EEG), magneto encephalography (MEG), positron emission tomography (PET), and fMRI. There are many limitations to the use of these techniques: they are expensive to operate and results are open to subjective interpretation. They are also intrusive for subjects although to different degrees. EEG and MEG are considered the least intrusive, while fMRI may cause many subjects discomfort related to having to lie still in a small space. Researchers must take the degree of intrusion into account since it affects data gathering, particularly if the subjects are busy managers and key decision-makers. Nevertheless, these techniques offer by far the best physical evidence to date on the activation of specific parts of the brain consequent to given stimuli.

Table 2
Example of How Constructs Could Be Defined and Measured in a Multi-Armed Gambling Task

Construct	Definition	Possible measure
Situation uncertainty	It can be defined in terms of the following: 1. Frequency (# times task is repeated during x time). 2. Heterogeneity (degree of novelty); 3. Causal ambiguity (number and degree of interdependence of subtasks; degree of simultaneity among subtasks) (Zollo and Winter, 2002)	Payoffs change randomly from trial to trial (manipulation of different situation characteristics)
Perceived utility	How much the individual likes the possible outcomes of the present situation (taking into account the anticipation of current outcomes and the memory of past decisions) (Aston-Jones and Cohen, 2005) (high−low on a scale depending on method)	ACC and OFC functioning using fMRI
Attention mode	What type of attention is devoted to a situation: broad or narrow (Aston-Jones and Cohen, 2005)	LC operation mode using fMRI or proxy with pupil diameter
Behavior	How much search for new alternatives is done (high being exploration, low exploitation) (Aston-Jones and Cohen 2005)	The observed strategy followed by the individual (i.e. her/his choice in each trial)

Note. ACC = anterior cingulate cortex; OFC = orbitofrontal cortex; fMRI = functional Magnetic Resonance Imaging; LC = locus coeruleus.

For the purpose of testing the ideas proposed in this manuscript,[3] the first best option is certainly to correlate the direct measures of the managers' attention focus in different real-life situations with the performance obtained in the real decisions (both at a purely individual level and at an organizational level) both in short- and long-term. However, given the difficulty (or impossibility) of measuring managers in real situations with the actual brain imaging tools, a good viable proxy would be to correlate attention focus measures obtained while performing different neuropsychological tasks in a lab context, with for example, managerial decision-making simulations facing the exploration–exploitation dilemma. The performance measures obtained in these simulations can be compared with the measures obtained in self-reported scales such as the cognitive flexibility scale (Martin & Rubin, 1995) or the recent scale of exploration−exploitation activities by Mom et al. (2009).

Conclusions

This paper addresses a dilemma common to organizations. Managing the trade-off between exploitation and exploration is fundamental to adaptive behavior and learning in increasingly complex and rapidly changing contexts. Although there has been much research on this trade-off, there are still several key gaps in the literature. First, we still know little about "how" exploration and exploitation are actually done. Second, the appropriate level of analysis at which the exploration−exploitation tradeoff is solved is not clear. Third, in many cases what is meant by exploration−exploitation is not very clear. Fourth, variation among individuals in the tendency to respond in an exploitative or explorative way to a given stimulus, and on the ability to shift their responses according to changes in the environmental conditions, have not been explored. To start addressing these key analytical gaps, we thus propose a framework and a method which, in our view, contributes particularly to the microlevel problem of individual-level variance in behavior. We provide a definition that is compatible with the management and the neuroscientific literature that we are using. We focus on the individual as the fundamental unit of analysis and study how exploration and exploitation are done at the micro, neurological level in terms of the processes going on in an individual's mind and the ensuing behavior when faced with a given environmental stimulus. Why does this framework matter? We believe there are at least four pos-

[3] While at the time of writing the study was still in its infancy, at the time of publishing the ideas were implemented along lines consistent with those here discussed and very promising preliminary results started to be obtained.

sible areas of contribution for the ideas put forward in this paper.

"How" the Tradeoff Is Solved—a Micro Perspective

As discussed earlier in this paper, there is still a remarkable lack of clarity as to what is the appropriate level of analysis for understanding the trade-off between exploitation and exploration. At the organizational level, many have focused on the attributes that make organizations more or less explorative (e.g., Tushman & O'Reilly, 1996). Others have argued instead (e.g., Papo, 2007) that individual-level processes ought to play a more central role in understanding the origin of new ideas. Obviously, these issues are related. Yet, we still know very little about how microlevel processes build up to organizational outcomes. We used the example of Menlo Park's activities to argue that the crucial issue here is not that of allocating exploitative versus explorative tasks to those individuals or organizations which are best suited to that task. Rather, our framework points to the idea that what matters to balance the tradeoff is not specialization, but the ability of key decision-makers to shift seamlessly from one task to the other, as we illustrated with the process Edison enacted in one of his key innovations. At the organizational level, this argument is consistent with recent research on the dynamics of innovation in complex technical systems, which has warned against the dangers of strategic outsourcing, often grounded in the belief that exploitation and exploration are actually separate activities which can be attributed to different types of organizations (e.g., Brusoni, Prencipe, & Pavitt, 2001). While there is some evidence of the dangers of outsourcing stemming from excessive specialization, microfounded explanations are in short supply. The application of neuroscience to managerial decision-making offers important opportunities for this line of inquiry to be developed to include the study of interactions among individuals engaged in these high-level cognitive tasks, and the generation of collective results that go beyond the sum of individuals' capacities. While the "aggregation" of individual capacities and behaviors into organizational ones remains a main limitation from the neuroeconomics field (Papo, 2007), this paper attempts, at least, to identify similarities between micro- and organization-level processes which ought to be further explored to generate sensible microfoundations of macrobehaviors.

"How" the Tradeoff Is Solved—a Macro Perspective

Domains that have traditionally focused on macro-organizational analyses such as organization theory, strategy, and entrepreneurship might benefit from the development of a theory of exploration and exploitation based on the neurological processes occurring within individual managers' brains.

For example, consider one of the basic insights in March's (1991) seminal article, which is related to the fact that exploration requires heterogeneity in human capital, which disappears without constant personnel turnover because the newcomers learn the code and adapt to the firm's modus operandi. March concludes (somewhat paradoxically) that the slower the learning of the code, the higher the exploration in the firm. However, the analysis in this paper points to an alternative mechanism which does not rely on personnel turnover to balance the exploration–exploitation dilemma, but rather focuses on the organizational members' neurological characteristics, that is, their propensity to function in tonic or phasic mode in the neuromodulation of their attention. It could be speculated that a broad attention span caused by the tonic activity of LC leads to a relatively slower socialization process (i.e., learning of the code) and thus more exploration at organizational level.

For scholars engaged in the study of entrepreneurship, our work offers the possibility to theorize and empirically validate (with objective measures) insights related to the explanation of explorative decisions, which might generate the foundation of new ventures internal to the firm (entrepreneurship) or outside it (spinoffs or start-ups). The launch of a novel enterprise can be viewed as the consequence of the continuous allocation of attention to explorative processes, and the development of neurological foundations for such choices can be particularly useful to this field. For similar reasons, strategy scholars might be able to build on a better understanding of the neurological foundations of exploration–exploitation decisions to develop models of strategic choice based on the idiosyncratic characteristics of the strategy-makers, over and above the influence of the industry context and the organizational endowment.

Enhancing Attention Modulation

Third, and more practically, the framework we propose in this paper might be used to develop teaching tools to foster individuals' ability to control and properly shift the focus of their attention. There is an element of plasticity in the neurological processes we discuss here, which can be leveraged to develop training programs to improve individuals' responses to changing environmental circumstances. Since attention modulation affects the ability to properly make decisions regarding exploration–exploitation, it will be then possible to adapt some instruments which have been developed for improving the attentional control (Sohlberg & Mateer, 1989, 2001). Very interestingly, for example, the analysis of the modulation of attention and its impact on decision-making might open new venues in managerial and entrepreneurial education through the development of tools and techniques which management students and practitioners can use to enhance their attentional control. These tools and techniques can be very varied, ranging from traditional in-class methods (e.g., for acquiring awareness of perceptive biases) to mental control practices (as, e.g., neuroimaging studies have demonstrated the positive impact of meditative practices on the improvement of functional up-regulation of brain regions affecting attention control. For a recent review see Rubia, 2009). In addition, since many psychiatric disorders of higher level cognition are thought to be due to deficits of attention (Posner & Petersen, 1990) researchers could adjust the tools used to treat patients of attention disorders to be applied to managers aiming to improve their attention control.

Reflection and Some Convergence Among Different Domains

On a broader note, a fourth contribution aims at reaching out to the now-broad crowd in social sciences interested in exploring the neuroscientific foundations of economic and organizational behavior. Social sciences are highly segmented in different knowledge domains. The findings of our work might foster a more integrated discussion. For example, mainstream economists might consider it interesting to move beyond the replication of results based on game theory, the focus of most neuroeconomics work so far, to develop novel insights into the neurological foundations of economic behavior. Also, the shift toward the explanation of variance across individual traits and consequent behavior, as opposed to the characterization of the mean tendency for the population of decision-makers, might serve to correct a standard bias in the discipline.

On the other side of the discipline's fence, evolutionary economists will appreciate our attempt to develop some of the microfoundations of the work related to how firms learn and evolve. The appropriate balancing of explorative and exploitative search is arguably a cornerstone of their theories. Depending on the attention mode, a certain way to decompose the problem will emerge and the search space will be defined accordingly. Search, therefore, is a consequence of the attention mode. If the attention focus is broad, for example, problems will be decomposed in broader modules, while a narrow attention focus will result in problems being decomposed in finer modules. Brusoni, Marengo, Prencipe, and Valente (2007) show that the balance to this trade-off depends upon the volatility of the problem environment. In stationary environments there is an evolutionary advantage to overmodularization, while in highly volatile environments, contrary to the received wisdom, modular search is inefficient in the long run. Similarly, we propose that in low-uncertainty tasks, exploitative behavior (i.e., deriving from a narrow attention mode and involving high decomposition of problems) would be advantageous, while the opposite will be true in the case of highly uncertain tasks: high decomposition is not efficient and explorative behavior derived from broad search will lead to better performance.

The organizational behavior field, with its orientation toward studying individual behavior in firms, is the obvious audience for the ideas discussed here. The main value for scholars in this field is that our findings expand their research agenda, which traditionally has been built on the application of social and cognitive psychology, to include the novel insights and methodologies developed in the neurosciences which are applicable to the study of organizational behavior phenomena. While neurologically founded findings might in the end confirm what many have already argued on the strength of sociometric surveys and psychological tests, neurosciences might contribute by offering better theoretical foundations to understand exactly why individuals behave in a certain way when isolated, and how organizational life is impacted by their responses.

Also, neuroscientists might find of interest the application of recent findings about the modulation of attention to the explanation of economic behavior and performance. The empirical validation of conceptual advancement in different decision-making settings and the search for differences among healthy people engaged in real-life, high-level decisions might offer good opportunities for the development of further understanding of the neural correlates of some of the highest functions of the human brain: the ability to search for novel approaches to the solution of unstructured problems in natural settings.

Future Research

Of course, many other things might impact individuals' ability to shift their focus of attention. The framework we have discussed here emphasizes the cognitive, rational elements of response to stimuli at a given point in time. Future research might explore other ways to approach the exploration–exploitation dilemma. For example, the balance between exploration and exploitation seems to be sensitive to time horizons, and humans show a greater tendency to explore when there is more time left for a task, presumably because it allows sufficient time to enjoy the fruits of those explorations (Carstensen, Isaacowitz, & Charles, 1999). Besides, individuals' impulsivity or differences in perception of time influence decision-making (Wittmann & Paulus, 2009), and so it would be useful to take the dimension of time into account as the exploration–exploitation dilemma encompasses the anticipation of future rewards. It would also be interesting to explore how emotions affect the management of the exploration–exploitation dilemma, since subjects whose brain activity displays good cooperation between the limbic area of the brain (emotional area) and the prefrontal cortex (thinking area) are the most successful in games based on experimentation (Bhatt & Camerer, 2005). Finally, future work could examine the tendency toward exploiting rather than balancing exploring−exploiting from an impatient behavior point of view, using the findings on the neural correlates of time discounting (McClure et al., 2004). As well as empirically testing the model proposed here, these ideas would provide a more complete understanding of the exploration−exploitation dilemma at the micro level.

To conclude, it should be noted that the data on brain functionality and the use of neuroimaging techniques have only begun to demonstrate their utility in complementing existing data sources, deepening the microfoundations of managerial behavior and developing better tools for improving cognitive abilities. The combination of novel empirical techniques and the objective strength of neuroscientific evidence, free from typical limitations of subjective data collection processes, to ground the theories of managerial behavior, demand increasing attention from management scholars. This becomes even more important if we include the potential to contribute to the ultimate goal of social science—as proposed by Glimcher and Rustichini (2004): the development of a unified theory of human behavior, without which the advancement of our understanding of managerial behavior cannot progress.

References

Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus–norepinephrine function: Adaptive gain and optimal performance. *Annual Review of Neuroscience, 28,* 403–450.

Aston-Jones, G., Rajkowski, J., Kubiak, P., & Alexinsky, T. (1994). Locus coeruleus neurons in the monkey are selectively activated by attended stimuli in a vigilance task. *Journal of Neuroscience, 14,* 4467–4480.

Barr, P. S., Stimpert, J. L., & Huff, A. S. (1992). Cognitive change, strategic action, and organizational renewal. *Strategic Management Journal, 13,* 15–36.

Bateman, T. S., & Zeithmal, C. P. (1989). The psychological context of strategic decisions: A model and convergent experimental findings. *Strategic Management Journal, 10,* 59–74.

Beckman, C. M. (2006). The influence of founding team company affiliations on firm behavior. *Academy of Management Journal, 49,* 741–758.

Beversdorf, D. Q., White, D. M., Chever, D. C., Hughes, J. D., & Bornstein, R. A. (2002). Central B-adrenergic modulation of cognitive flexibility. *Cognitive Neuroscience and Neuropsychology, 13,* 2505–2507.

Bhatt, M., & Camerer, C. F. (2005). Self-referential thinking and equilibrium as states of mind in games: FMRI evidence. *Games and Economic Behavior, 52,* 424–459.

Bodmer, U., & Vaughan, D. R. (2009). Approaches to preventing crises in family controlled small enterprises. *Journal of Neuroscience, Psychology, and Economics, 2,* 41–58.

Botvinick, M. M., Cohen, J. D., & Carter, C. S. (2004). Conflict monitoring and anterior cingulate cortex: An update. *Trends in Cognitive Sciences, 8,* 539–546.

Brusoni, S., Luigi, M., Prencipe, A., & Valente, M. (2007). The value and costs of modularity: A problem-solving perspective. *European Management Review, 4,* 15–39.

Brusoni, S., Prencipe, A., & Pavitt, K. (2001). Knowledge specialisation, organizational coupling and the boundaries of the firm: Why firms know more than they make? *Administrative Science Quarterly, 46,* 597–621.

Carlson, W. B. (2003). Invention and evolution: The case of Edison's sketches of the telephone. In J. Ziman (Ed.), *Technological innovation as an evolutionary process*. Cambridge, UK: Cambridge University Press.

Carstensen, L. L., Isaacowitz, D., & Charles, S. T. (1999). Taking time seriously: A theory of socioemotional selectivity. *American Psychology, 54,* 165–181.

Cohen, J. D., Aston-Jones, G., & Gilzenrat, M. S. (2004). A systems-level perspective on attention and cognitive control, guided activation, adaptive gating, conflict monitoring, and exploitation versus exploration. In M. I. Posner (Ed.), *Cognitive neuroscience of attention* (pp. 71–90). New York: Guilford Press Publications.

Cohen, J. D., McClure, S. M., & Yu, A. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society, 362,* 933–942.

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

D'Aveni, R. A., & McMillan, I. C. (1990). Crisis and the content of managerial communications: A study of the focus of top managers in surviving and failing firms. *Administrative Science Quarterly, 35,* 634–657.

Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature, 441,* 876–879.

Eisenberger, N. I., Lieberman, M. D., & Williams, K. D. (2003). Does rejection hurt? An fMRI study of social exclusion. *Science, 302,* 290–292.

Felin, T., & Foss, N. (2005). Strategic organization: A field in search of micro-foundations. *Strategic Organization, 3,* 441–455.

Felin, T., & Hesterly, W. S. (2007). The knowledge-based view, nested heterogeneity, and new value creation: Philosophical considerations on the locus of knowledge. *Academy of Management Review, 32,* 172–202.

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives, 19,* 25–42.

Gibson, C. B., & Birkinshaw, J. (2004). The antecedents, consequences and mediating role of organizational ambidexterity. *Academy of Management Journal, 47,* 209–226.

Gilzenrat, M. S., Cohen, J. D., Rajkowski, J., & Aston-Jones, G. (2003). Pupil dynamics predict changes in task engagement mediated by locus coeruleus. [Abstract] *Society for Neuroscience Abstracts No. 515.19.*

Glimcher, P. W., & Rustichini, A. (2004). Neuroeconomics: The concilience of brain and decision. *Science, 306,* 447–452.

Grève, H. R. (2003). *Organizational learning from performance feedback: A behavioral perspective on innovation and change*. Cambridge, UK: Cambridge University Press.

Gupta, A. K., Smith, K. G., & Shalley, C. E. (2006). The interplay between exploration and exploitation. *Academy of Management Journal, 49,* 693–706.

Hargadon, A., & Sutton, R. I. (2000). Building and innovation factory. *Harvard Business Review, May–June*, 157–166.

Holmqvist, M. (2004). Experiential learning processes of exploitation and exploration within and between organizations: An empirical study of product development. *Organization Science, 15,* 70–81.

Kaufman, J. C., & Baer, J. (2006). *Creativity and reason in cognitive development*. Cambridge, UK: Cambridge University Press.

Knutson, B., Fong, G. W., Bennett, S. M., Adams, C. M., & Hommer, D. (2003). A region of mesial prefrontal cortex tracks monetarily rewarding outcomes: Characterization with rapid event related fMRI. *NeuroImage, 18,* 263–272.

Koepp, M. J., Gunn, R. N., Lawrence, A. D., Cunningham, V. J., Dagher, A., Jones, T., . . . Grasby,

P. M. (1998). Evidence for striatal dopamine release during a video game. *Nature, 393,* 266–268.

Kolb, B., & Whishaw, I. Q. (1998). Brain plasticity and behavior. *Annual Review of Psychology, 49,* 43–54.

Kurke, L. B., & Aldrich, H. E. (1983). Mintzberg was right! A replication and extension of the nature of managerial work. *Management Science, 29,* 975–984.

Leonard-Barton, D. (1995). *Wellsprings of knowledge.* Cambridge, MA: Harvard Business School Press.

Levinthal, D. A., & March, J. G. (1993). The Myopia of learning. *Strategic Management Journal, 14,* 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.

Lubatkin, M. H., Simsek, Z., Ling, Y., & Veiga, J. F. (2006). Ambidexterity and performance in small-to medium-sized firms: The pivotal role of top management team behavioral integration. *Journal of Management, 32,* 646–672.

Lubinski, D., & Humphreys, L. (1997). Intelligence. *Incorporating General Intelligence into Epidemiology and the Social Sciences, 24,* 159–201.

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

March, J. G. (2006). Rationality, foolishness and adaptive intelligence. *Strategic Management Journal, 27,* 201–214.

March, J. G., & Simon, H. A. (1958). *Organizations.* New York: Wiley.

Martin, M. M., & Rubin, R. B. (1995). A new measure of cognitive flexibility. *Psychological Reports, 76,* 623–626.

McClure, S. M., Gilzenrat, M. S., & Cohen, J. D. (2006). An exploration–exploitation model based on norepinephrine and dopamine activity. In Y. Weiss, B. Sholkopf, & J. Platt (Eds.), *Advances in neural information processing systems* (Vol. 18, pp. 867–874). Cambridge, MA: MIT Press.

McClure, S. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science, 306,* 503–507.

McKenzie, J., Woolf, N., van Winkelen, C., & Morgan, C. (2009). Cognition in strategic decision-making: A model of nonconventional thinking capacities for complex situations. *Management Decision, 47,* 209–232.

Mintzberg, H. (1973). *The nature of managerial work.* New York: Harper & Row.

Mom, T. J. M., Van Den Bosch, F. A. J., & Volberda, H. W. (2007). Investigating managers' exploration and exploitation activities: the influence of top-down, bottom-up, and horizontal knowledge inflows. *Journal of Management Studies, 44,* 910–931.

Mom, T. J. M., Van Den Bosch, F. A. J., & Volberda, H. W. (2009). Understanding variation in managers' ambidexterity; investigating direct and interaction effects of formal structural and personal coordination mechanisms. *Organization Science, 20,* 812–828.

Montague, P. R., King-Casas, B., & Cohen, J. D. (2006). Imaging valuation models in human choice. *Annual Review of Neuroscience, 29,* 417–448.

Murray, E. A., O'Doherty, J. P., & Schoenbaum, G. (2007). What we know and do not know about the functions of the orbitofrontal cortex after 20 years of cross-species studies. *Journal of Neuroscience, 27,* 8166–8169.

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

O'Reilly, C. A., & Tushman, M. (2004, April). The ambidextrous organization. *Harvard Business Review,* 1–8.

Papo, D. (2007). Where do ideas come from? Brain activity and economic systems. *European Journal of Economic and Social Systems, 20,* 163–186.

Perretti, F., & Negro, G. (2006). Filling empty seats: How status and organizational hierarchies affect exploration and exploitation in team design. *Academy of Management Journal, 49,* 759–777.

Peyron, R., Laurent, B., & García-Larrea, L. (2000). Functional imaging of brain responses to pain. A review and meta-analysis (2000). *Neurophysiologie Clinique/Clinical Neurophysiology, 30,* 263–288.

Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. *Annual Review of Neuroscience, 13,* 25–42.

Raisch, S., & Birkinshaw, J. (2008). Organizational ambidexterity: Antecedents, outcomes, and moderators. *Journal of Management, 34,* 375–409.

Rolls, E. (2000). *The brain and emotion.* Oxford, UK: Oxford University Press.

Rothaermel, F. T., & Hess, A. M. (2007). Building dynamic capabilities: Innovation driven by individual-, firm-, and network-level effects. *Organization Science, 18,* 898–921.

Rubia, K. (2009). The neurobiology of Meditation and its clinical effectiveness in psychiatric disorders. *Biological Psychology, 82,* 1–11.

Sara, S. J. (2009). The locus coeruleus and noradrenergic modulation of cognition. *Nature Reviews Neuroscience, 10,* 211–223.

Shinoda, T., & Kato, M. (2006, October). *A pupil diameter measurement system for accident prevention systems, man and cybernetics.* Paper presented at the IEEE International Conference, Taipei, Taiwan.

Simon, H. A. (1985). Human nature in politics. *American Political Science Review, 79,* 293–304.
Simon, M., Houghton, S. M., & Aquino, K. (2000). Cognitive biases, risk perception, and venture formation: How individuals decide to start companies. *Journal of Business Venturing, 15,* 113–134.
Skrabec, Q. R. (2006). *George Westinghouse: Gentle genius*. New York: Algora Publishing.
Smith, W. K., & Tushman, M. (2005). Managing strategic contradictions: A top management model for managing innovation streams. *Organization Science, 16*(5), 522–536.
Sohlberg, M., & Mateer, C. (1989). *Introduction to cognitive rehabilitation: Theory and practice.* New York: Guilford Press.
Sohlberg, M., & Mateer, C. (2001). *Cognitive rehabilitation: An integrative neuropsychological approach.* New York: Guilford Press.
Sterpenich, V., D'Argembeau, A., Desseilles, M., Balteau, E., Albouy, G., Vandewalle, G., et al. (2006). The locus ceruleus is involved in the successful retrieval of emotional memories in humans. *The Journal of Neuroscience, 26,* 7416–7423.
Swedberg, R. (1993). *Explorations in economic sociology*. New York: Russell Sage Foundation Publications.
Taub, E. (2004). Harnessing brain plasticity through behavioral techniques to produce new treatments in neurorehabilitation. *American Psychologist, 59,* 692–704.
Thomas, J. B., Clark, S. M., & Gioia, D. (1993). Strategy sensemaking and organizational performance: Linkages among scanning, interpretation, action, and outcomes. *Academy of Management Journal, 36,* 239–270.
Tom, S. M., Fox, C. R., Trepel, C., & Poldrack, R. A. (2007). The neural basis of loss aversion in decision-making under risk. *Science, 315,* 515–518.
Tushman, M., & O'Reilly, C. A. (1996). Ambidextrous organizations: Managing evolutionary and revolutionary change. *California Management Review, 38,* 8–30.
Usher, M., Cohen, J. D., Servan-Schreiber, D., Rajkowski, J., & Aston-Jones, G. (1999). The role of the locus coeruleus in the regulation of cognitive performance. *Science, 283,* 549–554.
Volberda, H., Baden-Fuller, C., & Van Den Bosch, F. A. J. (2001). Mastering strategic renewal: Mobilizing renewal journeys in multi-unit firms. *Long Range Planning, 34,* 159–178.
Weick, K. E., & Sutcliffe, K. M. (2006). Mindfulness and the quality of organizational attention. *Organization Science, 17,* 514–524.
Wittmann, M., & Paulus, M. P. (2009). Temporal horizons in decision making. *Journal of Neuroscience, Psychology, and Economics, 2,* 1–11.
Yu, A., & Dayan, P. (2005). Uncertainty, neuromodulation and attention. *Neuron, 46,* 681–692.
Zollo, M., & Winter, S. G. (2002). Deliberate learning and the evolution of organizational capabilities. *Organization Science, 13,* 339–352.