M@n@gement
2018, vol. 21(3): 1050-1079

The Continuum Conception of Exploration and Exploitation: An Update to March's Theory

Sasanka Sekhar Chanda ● Sougata Ray ● Bill Mckelvey

Accepted by co-editor in chief Thomas Roulet

Abstract. In his seminal 1991 publication, March illustrates the continuum conception of exploration and exploitation by an organizational learning metaphor. Exploration involves allocating resources to experimentation. Exploitation involves doing known things better and focusing on execution. In March's formal model, members of an organization deploy their collective human capital and engage in learning activities to fashion organizational knowledge. Collective human capital (CHC) is constituted by the aggregate beliefs of members, some of which are correctly aligned with respect to an objective external reality while others are neutral or misaligned. Organizational knowledge—constituting the validated knowledge in an organization—resides in the databases, rules, forms, norms, operating procedures and other artifacts in an organization. March's computational experiments pertaining to the continuum conception suggest that more exploration is always preferable over more exploitation. We demonstrate that the reverse holds true when the CHC available in an organization is somewhat lower than that assumed in March's experiments. Our research indicates that a section of extant research is mistaken in assuming that March's formal model for the continuum conception suggests an inverted U-shaped relation between the extent of exploration and organizational outcome. Instead, the level of CHC determines whether it is rewarding to focus on exploration or exploitation. Thus, the formal model supports managerial intentionality towards exploratory and exploitative innovation through appropriate choice of the level of CHC. We call for a new "balance" discussion, focusing on the determinants of the minimum level of the non-preferred activity from among exploration and exploitation.

Keywords: bottom of the pyramid, computational simulation, exploration, exploitation, human capital, organizational knowledge, organizational learning

Sasanka Sekhar Chanda
Indian Institute of Management Indore
India
sschanda@iimidr.ac.in

Sougata Ray
Indian Institute of Management Calcutta
India
sougata@iimcal.ac.in

Bill Mckelvey
Anderson School of Management
Los Angeles, CA
mckelvey@anderson.ucla.edu

> *"I have come to believe that […] ideas that transform ways of thinking about practical problems rarely come from a direct focus on those problems."* March (2013: 732)

INTRODUCTION

Organizations are advised to strike a balance between exploration (allocating resources to experimentation) and exploitation (focus on execution and doing known things better). Overemphasis on exploitation will trap an organization in sub-optimal equilibria and result in obsolescence of an organization's outputs. Overemphasis on exploration fails to generate a sufficient number of timely wins[1] and may lead to an organization losing its way in its perpetual exploration of novelty (Levitt & March, 1988; Schmitt, Probst & Tushman, 2010).

Extant literature is unable to provide a satisfactory answer to the question as to what mix of exploration–exploitation rates constitute an appropriate balance. In a stable environment—characterized by little change to environmental parameters salient to an organization—March (1991: 77, Figure 2) appears to suggest that higher exploration is always desirable over higher exploitation. In this formalization, exploration and exploitation are considered as two ends of a continuum. One is at a loss to figure how this state of affairs brings a balance to organizational functioning. In our study, we seek to address this puzzle.

Notwithstanding the arguments underscoring the necessity of balance between exploration and exploitation presented in March's text, it is clear that the implications of March (1991: 77, Figure 2)[2] have the makings of an incomplete prescription. March's study has inspired generations of researchers. As of March 2018, over 20,000 citations to the study can be found in Google Scholar. Yet, until and unless we identify a situation where higher exploration is *not* desirable over lower exploration, we have an anomaly: arguments grounded in organizational reality suggest the necessity of balance; the published results from the formal model do not. The anomaly should not stay unresolved, since there are important consequences for theory and practice. Managers need to decide how many resources should be apportioned to exploration activities and how many resources should be allocated to exploitation. Managers seek guidance as to what resource allocation mix is more desirable in the various stages of organizational life (Gupta, Smith & Shalley, 2006). Resolution of the anomaly will help to fulfill managerial expectations from scholarship.

Posen and Levinthal (2012) attempt to answer the question implicit above: *what mixes of resource allocation to exploration and exploitation are desirable in stable and dynamic environments?* Posen and Levinthal (2012) use a single-actor, multi-armed bandit model to show that: (a) a 50:50 mix is desirable in stable environments (p. 592, Figure 1) and (b) when the rate of environmental change increases, a higher proportion of exploitation is desirable (p. 594, Figure 3). An important factor motivating our study is the curiosity to know whether the model in the study originating the concepts of exploration and exploitation in organization and management studies—March (1991)—would provide the same answer.

In our study, we recreate March's computational model, based on the conceptual description given in his 1991 article. An agent-based computational model is used because it is impossible to experiment with a real organization by putting it into an experimental-laboratory situation to test what philosophers of science call a "counterfactual conditional", i.e. truly test the following: If *A* exists, *B* will exist; if *A* is removed, *B* will disappear (Fillenbaum 1974; Pearl 2000).

1. By "timely" we imply that knowledge acquired by exploration needs to be converted into marketable products—an endeavour that requires exploitation—within reasonable time frames, so as to ensure survival and continued relevance of the organization.

2. The continuum conception of exploration and exploitation is illustrated by the curve "HETEROGENEOUS p1" in March (1991: 77, Figure 2).

March's study draws from the work on genetic algorithms by Holland (1975). Members of an organization begin with a random set of beliefs. The organization's members learn *only* from the brain of the organization, i.e. its organizational code (hereafter referred to as org. code).[3] The latter is a repository of *organizational knowledge* (Gavetti, Greve, Levinthal & Ocasio, 2012). Organizational knowledge resides in the databases, rules, forms, norms, operating procedures and other artifacts in an organization (March, 1991; Nelson & Winter, 1982). The org. code, in turn, learns from the elites among organizational members, i.e. from those members who are more knowledgeable about the external reality or the relevant environment of an organization than the code itself. Equilibrium is reached when the set of correct beliefs of members converges to the set of dimensions on which the org. code has correct knowledge (i.e. values conforming to that in the corresponding dimensions in the external reality).

We extend the model to allow variation of the *collective human capital* (CHC) (von Nordenflycht, 2011) in the organization. Human capital consists of "…the knowledge, skills, and abilities a firm's workers acquire as a result of their learning, education, training, and experience" (Becker, 1964; quoted in Lee, Bachrach & Rousseau, 2015: 796). Human capital is frequently considered as a property of an individual. In this study, we invoke von Nordenflycht's characterization of CHC to refer to the aggregate human capital residing in organizational members. In March's model, members of the organization—the "agents" in the model—were endowed with random values in their knowledge dimensions at initialization. Equal proportions of the beliefs of members were aligned, misaligned or neutral with respect to the environmental requirements. The CHC so endowed can be reduced or enhanced by increasing the proportion of misaligned beliefs and by increasing the proportion of aligned beliefs, respectively.

We conduct computational simulation experiments to observe the level of organizational knowledge attained when CHC is varied. We find that when an organization initiates a task for which it possesses grossly deficient CHC, a higher level of organizational knowledge results when the rate of exploitation is higher compared to a case where the rate of exploration is higher (i.e. exploitation rate is lower). Conversely, if the CHC is at the level of that endowed by March in his experiments or higher, high values of organizational knowledge are attained when exploration is high—i.e. when the rate of exploitation is low. We also find that the level of organizational knowledge attainable from moderate or high CHC is more than twice that attainable from low CHC. Thus, not only are the strategies for success different on either side of a critical mass of CHC, the levels of organizational attainments are vastly different as well.

Further investigation informs us about the mechanism by which the switch from desirability of higher exploration to desirability of higher exploitation materializes. The findings from the formal model refute a dominant belief—observed in extant research—that there is an inverted U-shaped relationship between extent of exploration and organizational outcomes. Rather, the level of endowment of CHC determines whether an organization shall find it remunerative to focus on exploration or exploitation. Nevertheless, an organization must allocate a certain minimum level of resources to the activity it chooses not to focus on. This suggests that a new discussion on exploration–exploitation balance is necessary, centered on the following question: *Given the enabling of one from exploratory or exploitative innovation via the level of CHC of an organization, what are the determinants of the minimum levels of the other*

3. By design, the effect of any learning of one member from another is considered negligible.

activity, in specific contexts?

In the section that follows, we briefly review extant literature on exploration and exploitation. Then we discuss the theoretical foundations of the study. Thereafter we derive some propositions for testing by means of the formal model. The simulation model is described next, followed by the results in graphical form and our interpretation thereof. We discuss the significance of the results with respect to examples from practice. We end by highlighting the implications and limitations of our study, and suggest some avenues for future research.

LITERATURE REVIEW

The literature on exploration and exploitation following March progresses along two principal tracks. One track comprises computational simulation studies that seeks to build new theory by extensions to March's model. The second track comprises construction of theory without the reliance on a formal (mathematical or computational or other formal logic) model[4], extending March (1991), and testing of such theory by empirical studies.

The track comprising computational simulation studies may be further categorized into studies that use a part or whole of March's formal model and studies that use a formal model with several assumptions different from March. In the former category we find the following studies: Blaschke and Schoeneborn (2006), Chanda (2017), Chanda and Ray (2015), Fang, Lee and Schilling (2010), Kane and Alavi (2007), Kim and Rhee (2009), Miller and Martignoni (2016), Miller, Zhao and Calantone (2006), Rodan (2005), Schilling and Fang (2014) and Zhang and Xi (2010).[5] The studies by Lazer and Friedman (2007), Posen and Levinthal (2012) and Siggelkow and Rivkin (2006) illustrate addressing exploration–exploitation themes by using a formal model different from March (1991). In Table 1 we present the findings of the studies in this track in brief. The investigation—with the set of all modeling studies drawing from March (1991)—assures us that the work we carry out has not already been reported in prior research.

4. Theory without reliance on a formal (mathematical or computational or other formal logic) model is sometimes referred to as *verbal theory* or *informal theory* (for example Bendor, Moe & Shotts, 2001). We refer to such kind of theory as *verbal theory* in the rest of the paper. See also, Adner, Pólos, Ryall & Sorenson (2009).

5. Unfortunately, we were unable to access Zhang and Xi (2010). From references to their paper in other studies, we learn that Zhang and Xi (2010) experimented with alternative heuristics for agents' selection of partners from whom to learn.

Type of formal model	Description
Models similar to March's (1991) models	Rodan (2005), introduces a new construct, "self-restrained experimentation", by randomly assigning ("+1" / "-1") values to neutral ("0") beliefs of members, at the time of selection of elite members from whom the organizational code learns. A key finding is that, in changing environments, inclusion in the policy-making elite is more effective when predicated on near-term performance, than on tenure. Stringent selection criteria are desirable. Blaschke and Schoeneborn (2006) model a new construct "forgetting by an individual", whereby an individual holding non-conforming beliefs passes on to a stage of holding a neutral belief, upon encountering a different belief when learning from the organizational code (hereafter referred to as org. code). Blaschke and Schoeneborn (2006) find that forgetting serves as a source of dynamic instability for an organization. Miller, et al. (2006) relax a key assumption of the genetic algorithm in March (1991) that members do not learn directly from each other. Rather, some knowledge is considered as tacit. Such knowledge would get transmitted only through direct interaction between organizational members (along with non-tacit or codifiable knowledge). Non-tacit knowledge would also get transmitted via the intermediation of the org. code, as in March (1991). They find that "the small-world effect of learning through distant search becomes redundant if the org. code facilitates knowledge transfer among distant individuals." (p. 716) Kane and Alavi (2007), inspect the effect of information technology-enabled learning mechanisms—email, knowledge repositories of best practices and groupware— on exploration and exploitation. Organizational members learn from each other as well as through information technology systems in the company. They find that, each of these IT-enabled learning mechanisms enable capabilities that have a distinctive effect on learning dynamics in the organization. Kim and Rhee (2009), permit learning between pairs of agents, in addition to the arrangement that agents learn from the org. code. They also varied the amplitude and frequency of environmental turbulence. Their research suggests that adaptations of organizational knowledge are facilitated better when internal variety is managed through a combination of strong complementary practices. Fang, et al. (2010), introduced variation in the network structure for interpersonal learning. Agents are situated in semi-isolated groups within an organization. Their experiments manipulate the number of ties between and within groups. There is no org. code to mediate learning, unlike the March (1991) model. Fang, et al. (2010) find that moderate levels of cross-group linking lead to highest equilibrium performance. Schilling and Fang (2014), examined the role of "hub" individuals in the diffusion of learning in organizational social networks. Hub individuals have a disproportionately large number of ties, compared to other individuals. Mediation by the org. code is absent in this model. Schilling and Fang (2014) find that a network with a moderate number of hubs outperform networks with a high number of hubs and networks where hubs are rare. Further, information distortion at the hubs yields positive outcomes under some conditions. Chanda and Ray (2015) inspect the full state space of variation of exploration and exploitation, for the orthogonal conception described in March (1991). Chanda and Ray (2015) show that, there exists multiple exploration–exploitation combinations for which the outcome—organizational knowledge—is optimal (highest). They infer that presence of multiple optima—some emphasizing exploitation, others emphasizing exploration—justifies empirical instances of managers shaping their organizations to be strongly into exploratory or exploitative innovation. Miller and Martignoni (2016) model interpersonal learning among agents in an organization, departing from March's approach comprising agents learning through intermediation of an org. code. They find that, when agents are subject to forgetting, a higher rate of interpersonal learning often enhances the diversity of beliefs within an organization. Chanda (2017) shows how managers' contributions towards organizational success may be evaluated despite that fact that circumstances beyond managers' control may impact outcomes. This is accomplished by modeling complexity along the lines of *Shannon Complexity* (also used in Duncan 1972 and Child 1972) as the minimum extent of org. code knowledge that must be configured correctly, in order that a firm succeeds.

Other computational simulation studies*	Siggelkow and Rivkin (2006) take recourse to an *NK* model to show that (a) extensive exploration at lower levels of a hierarchical organization leads to better performance when there are no cross-departmental interdependencies and (b) when interdependencies cut across the domains of low-level managers, higher exploration at low levels reduce overall exploration in environments that require broad search. Lazer and Friedman (2007) inspect how the structure of communication networks among actors affect system-level performance. They use an *NK* model to simulate different kinds of network structures in which agents are located. Lazer and Friedman (2007) find that moderately connected networks outperform poorly-connected and well-connected networks. Further, when agents are dealing with a complex problem, the more efficient the network at disseminating information, the better the short-run performance; but, at the same time, the long-run performance of the system is lower. Posen and Levinthal (2012) take recourse to a *multi-arm bandit model*. They consider the decision-maker as one facing a set of slot machines. Each slot machine has an unknown probability of yielding a certain payoff. Initially the slot machines are operated randomly for a few times. The probability of getting payoff are estimated. Thereafter, the exploration–exploitation dilemma is framed as follows: should one exploit prior knowledge, by choosing slot machines known to produce higher payoff, as from historical data; or should one try out a less-used slot machine, in the hope of discovering one with even higher payoff than the incumbents (i.e. slot machines whose payoff probabilities are better known). Posen and Levinthal (2012) find that a 50:50 distribution is optimal, in a stable environment (p. 592, Figure 1); for dynamic environment, the proportion changes, eventually favoring higher exploitation under higher turbulence (p. 594, Figure 3).

**This is an illustrative list, i.e. is not an exhaustive list of all studies that refer to exploration and exploitation and perform computational simulations.*

Table 1. Computational simulation studies referring to exploration and exploitation

The track comprising construction and testing of verbal theory referring to exploration and exploitation constructs is quite vast. A dominant theme is that of *ambidexterity*, i.e. the need for companies to make arrangements for carrying out both exploration and exploitation activities in the organization. In this context, Schmitt, et al. (2010: 129) suggest that although exploration and exploitation "… may be interrelated, they require underlying organizational processes, structures, strategies, and cultures that differ substantially. The ability to manage these conflicting demands is fundamental for sustainable performance". An important debate concerns whether: (a) exploration and exploitation should be housed in separate units that are loosely connected (*spatial ambidexterity*), for example, Tushman and O'Reilly (1996) and Benner and Tushman (2003); or (b) whether companies should pay sequential attention to exploration and exploitation (*temporal ambidexterity*), for example, Eisenhardt and Brown (1997) and Hamel and Prahalad (1993);[6] or (c) whether exploration and exploitation may be simultaneously operational in the same organizational unit at the same time (*contextual ambidexterity*), e.g. Gibson and Birkinshaw (2004).

It may be noted that papers by Adler, Goldoftas and Levine (1999) and Birkinshaw, Zimmerman and Raisch (2016) suggest that all three modes of ambidexterity are observable. This contributes to extinguishing the polemics of the debate. Further, in a critical vein, Lavie, Kang and Rosenkopf (2011) aver that both *spatial* and *temporal separation* involve challenges to management that may eclipse gains from balancing exploration and exploitation. In the former (spatial separation), senior managers have to integrate disparate output coming from the units doing either exploration or exploitation exclusively. In the latter (temporal

6. Hamel and Prahalad (1993) provide the example of *Komatsu*. At one time, *Komatsu* lagged its more well-known competitor, *Caterpillar*, in quality. *Komatsu* first initiated a total quality program, followed by value engineering and manufacturing rationalization (exploitation). Thereafter it focused on new product development (exploration). Having achieved a certain portfolio, it focused on increasing the speed of product development to attain variety at low cost (exploitation). Expansion to global markets intensified thereafter (exploration).

separation), senior managers have to repeatedly re-configure the organization whenever a switch happens. However, other research (for example, Brady & Davies, 2004; Garcias, Dalmasso & Sardas, 2015) suggests that the latter is not as big a problem as Lavie, et al. (2011) make it out to be. Managers may take recourse to transition learning (or exploitative learning), in which a company uses "capabilities located locally in an exploratory project to diffuse, replicate and extend such learning to other projects or to the organization as a whole" (Garcias, et al., 2015: 159).

Using verbal theory referring to exploration and exploitation, researchers have also investigated whether strategic alliances serve the need to carry out both exploration and exploitation. Some literature focuses on finding out which of exploration and/or exploitation is better done in-house or in the alliance, in specific contexts (e.g. Stettner & Lavie, 2014). Lavie, et al. (2011) suggest that organizations may carry out exploration and exploitation in different *domains*: organizations may opt for technology versus marketing and production alliances for the *function domain*, and do so with new or existing partners (*structure domain*).

We acknowledge that models like March's can be used to develop formal theory and enrich some of the debates referred to above, through formal comparison of relative merits. However, that can happen when two conditions are met: (a) it is possible to replicate March's model and (b) its inner workings are better known. Our paper seeks to contribute in these areas.

THEORETICAL FOUNDATIONS

An organization's survival and prosperity are strongly dependent on its ability to excel in exploration of new possibilities and exploitation of old certainties (Holland, 1975; Kuran, 1988; March, 1991; Schumpeter, 1934). Exploring new opportunities is future-looking and involves variance-increasing activities, say through experimentation (March, 1991; Schmitt, et al., 2010). Exploiting existing products and services is rooted in the past and involves variance-reducing activities (Farjoun, 2010; Smith, Binns & Tushman, 2010) via execution focus and refinement.

March elaborates on mechanisms underlying exploration and exploitation via an organizational learning metaphor. March demonstrates that organizations can improve their chances of attaining higher knowledge levels if they cultivate heterogeneous knowledge. The continuum mechanism involves a portion of the organizational members undergoing slower socialization to the ways of the organization. When a section of organizational members undergo slower socialization (i.e. learn slowly), the heterogeneity of their knowledge is preserved longer. Exploitation increases as average member learning rates increase, with a concomitant decrease in exploration. Thereby, in the continuum formalization, a higher proportion of *Fast Learners* relative to *Slow Learners* orients an organization towards higher exploitation and lower exploration; a lower proportion of *Fast Learners* relative to *Slow Learners* orients the organization towards higher exploration and lower exploitation.

March's research is a rare but highly influential instance of multilevel theory development that establishes a connection between the organizational-learning literature (e.g. Huber, 1991, Levinthal & March, 1993; Simon, 1991) and the knowledge-management literature (e.g. Hargadon & Fanelli, 2002; Nonaka & von Krogh, 2009; Spender, 2008). At the core of the formal model, *organizational knowledge* is generated by deployment of *collective human capital* of organizational members through stylized learning processes.[7] Simon (1991: 126, italics in original) observed that it is "…important to specify *where* in the organization particular knowledge is stored, or *who* has learned it". Accordingly, in the paragraphs that follow, we first elaborate on the two terms *collective human capital* and *organizational knowledge.* Thereafter, by means of a schematic diagram, we explain how March's formal model interconnects them by stylized learning processes.

At the level of the individual, knowledge constitutes an understanding of principles, facts and processes that range from generic to specific, e.g. knowledge of accounting to knowledge of how to use a particular company's accounting software (Ployhart & Moliterno, 2011). Hargadon and Fanelli (2002: 294) refer to stocks of knowledge held by individuals, as *latent knowledge*. In their view, "…the schemata—comprising scripts, goals, and identities—of members of an organization make up the latent knowledge available within that organization". This knowledge is dispersed in an organization in the brains of its employees. Some examples are: knowledge of peculiarities of sales territories (say, credit needs of customers and variation in retail stocking cycles); actual capacities and quirks of production machinery; issues involved in buying raw materials domestically vs. sourcing internationally, and so forth. For the purposes of this study, we say that the latent knowledge in the members of an organization is manifested as CHC, in the formal (computational) model in this study. We make a distinction between latent knowledge and CHC in order to distinguish the set of beliefs relevant to an organization (on which the organization fashions its activities and end-products) from the whole gamut of all possible matters (organizational or otherwise) that constitute latent knowledge in an individual.

Organizational knowledge is another construct pertaining to knowledge found in an organization. Nonaka and Takeuchi (1995) define organizational knowledge as the validated understanding and beliefs in an organization about the relationship between the organization and its environment. Examples may be found in the recipes a company puts into use: standard operating procedures; product manufacturing know-how consisting of technologies used, productive machinery deployed and quality standards espoused; distribution channel know-how comprising techniques for channel development, avoiding channel conflict, etc.; raw material sourcing know-how about vendor development, global sourcing, and so on. Even the structure and configuration of information in a company's databases and the roles and access control processes in information technology systems are manifestations of organizational knowledge.

Organizational knowledge creation is dependent on the ability of organization members to exchange and combine existing information, knowledge and ideas (Kogut & Zander, 1992, 1996; Smith, Collins & Clark, 2005). Significantly, though CHC is present in a dispersed form in the organization, it may not be used to serve the purposes of the organization unless assimilated into organizational knowledge. Organizational knowledge is the validated knowledge that is actually deployed for productive ends. It follows that the level of organizational knowledge that is attained is a function of the CHC in the company.

In Figure 1, we present a schematic of March's (1991) model that is within the scope of our study.[8] We identify two distinct knowledge stocks, *collective human capital* (*CHC*) and *organizational knowledge*, and the

7. We describe these learning processes after elaborating on the terms *collective human capital* and *organizational knowledge.*

8. March (1991) developed additional findings by considering the organization as an open system, where he allowed inflow of heterogeneous knowledge from outside the organization at a variable rate (pp. 79–80, Figures 4 and 5). This part of March's work is out of scope of our present study. Interested readers may please refer to Chanda and Ray (2015) and Chanda (2017).

hosts of such stock. As stated earlier, the CHC available to the organization (von Nordenflycht, 2011) is held by members of the organization.[9] Likewise, technical systems, databases, routines and procedures in use in the organization contain organizational knowledge (Levitt & March, 1988; Nelson & Winter, 1982). A third entity, the organization's competitive environment, shows the reality relevant to the task at hand. In our study the external reality is randomly initiated for each experiment, and remains constant throughout each experiment (i.e. model "run" through all the "time-steps"). This creates a stable environment for each run of the model.[10] The level of any knowledge stock is measured by the degree of conformance to the corresponding elements in the external reality.

In all, there are two flow variables, p_2 and p_1, which are defined in Figure 1. Both signify rate of learning. They differ in who learns from whom.

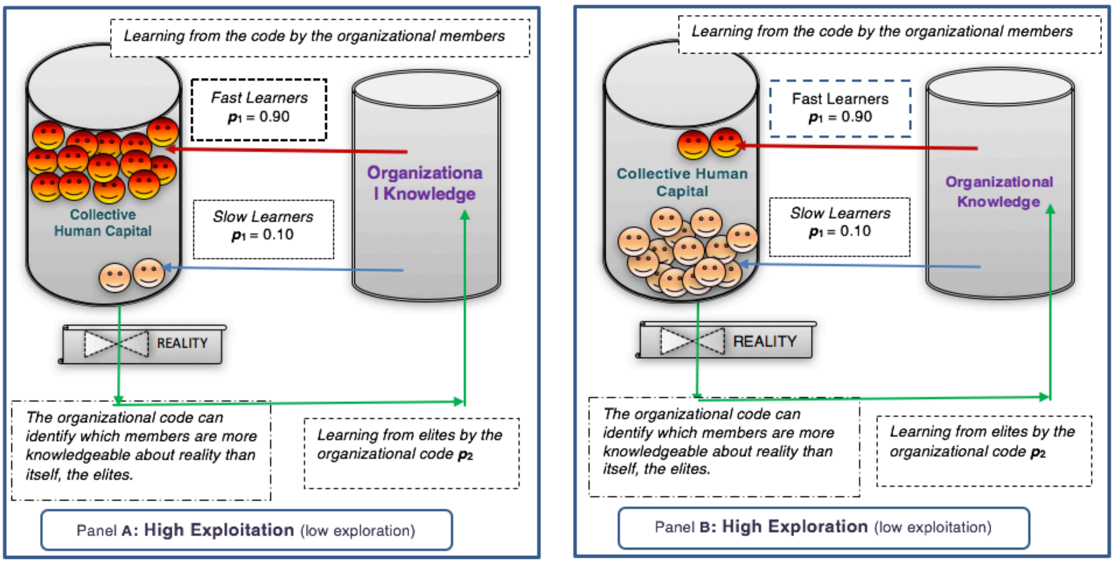

Notes. The stock of organizational knowledge is enhanced by means of the flow represented by the org. code learning at a rate p_2. The stock of collective human capital in the organization gets updated when individual members learn at rate p_1 (a flow variable). High exploitation occurs when p_{1-AVG} = 0.80 (Panel A). In this case, about one-eighth of the organizational members are *Slow Learners* and the rest are *Fast Learners*. High exploration occurs when p_{1-AVG} = 0.20 (Panel B). For this situation, about seven-eighths of organizational members are *Slow Learners* and the rest are *Fast Learners.* The continuum from high exploitation to high exploration is fashioned by decreasing proportion of *Fast Learners* relative to *Slow Learners*. Values in the external reality or environment determine the correctness or otherwise, of any knowledge dimension.

Figure 1. Stocks and flows in the model

No beliefs are defined in the initial org. code. The probability that the org. code updates itself with the value from a specific knowledge dimension of elite members (who know more about the external reality or relevant environment, compared to the org. code) is higher when p_2 is high and lower when p_2 is low.

9. We assume that organizational members are otherwise motivated to deploy the CHC for organizational purposes.

10. In the last four sets of experiments that we report (Figures 6, 7, 8 and 9), we relax this assumption to observe outcomes over a range of environmental turbulence.

Organization members that are *Slow Learners* learn from the org. code at a low rate (10%, p_1 = 0.1, as from March's study). This signifies that, in any given time-step, there is a low probability (10% chance) that a member will update his/her belief with that from org. code when a non-null value exists for a given dimension of knowledge in the org. code. Organization members that are *Fast Learners* learn from the org. code at a high rate (90%, p_1 = 0.9, following March). In March's formalization of exploration and exploitation as two ends of a continuum, a high rate of exploration is motivated by having a high proportion (about seven-eighths) of organizational members learn slowly. A high rate of exploitation is motivated by having a high proportion (about seven-eighths) of organizational members learning fast.

We note that validated knowledge accumulates in March's org. code by drawing from the CHC of the members. We define a *low* level of CHC as one where a significant proportion of the beliefs of organizational members are inaccurate representations of the reality corresponding to the task at hand. We define a *moderate* level of CHC as one where, on average, equal proportions of beliefs of members are aligned, misaligned or neutral. Populations having a high proportion of member beliefs aligned with the reality corresponding to the task at hand are deemed to have *high* CHC.

THEORY DEVELOPMENT

In the experiments reported by March (1991), the organizational population is given a similar amount of human capital at the beginning of all experiments. On average, the CHC has no significant misalignment with the needs of the task at hand, i.e. the results correspond to moderate levels of CHC.[11] March demonstrated that if individuals learn too fast, the organizational code knowledge reaches sub-optimal values, since differences between the code knowledge and knowledge of individuals quickly gets eroded. However, if individuals learn slower higher organizational knowledge is attained, since diversity is preserved longer. In other words, when a large fraction of organizational members learn slowly, knowledge heterogeneity is preserved for a longer period of time. Such a condition pertains to a state of high exploration. The other end of the continuum—where *Fast Learners* are in higher proportion—indicates a state of high exploitation. In the latter case, knowledge heterogeneity erodes more quickly, leading to lower organizational knowledge. This suggests the following proposition:

> **P1**: In organizations with a moderate level of CHC, the level of organizational knowledge attained from high exploration is greater than the level of organizational knowledge attained from high exploitation.

We further note that, absent any external influence, the CHC is the only reservoir that organizational knowledge can draw from. An organization begins with a low level of useful CHC if the beliefs of most organizational members are misaligned with respect to the demands imposed by its external competitive environment. In such cases, the level of organizational knowledge attained by learning activities will be capped to the extent that scarce environmentally relevant beliefs are rounded up and

11. We elaborate on this point when we describe the formal model in the next section. For the moment, it suffices to know that on an average, one-third of beliefs of each member comprise correct beliefs and one-third comprise misaligned or incorrect beliefs. The remaining third of beliefs of members are neutral beliefs (i.e. in the nature of "no opinion").

assimilated into the org. code. This level will be lower than the level of organizational knowledge attained when the CHC is at moderate or high levels, as is applicable in situations when an organization operates in a competitive environment in which its capabilities are well-suited to apply. This is because, in the latter case, the CHC in the organization is comprised of a higher proportion of correct or well-aligned beliefs held by its members. This suggests the following proposition:

> **P2**: Organizations with a low level of CHC develop lower organizational knowledge from exploitation and exploration activities, compared to organizations with a moderate or high level of CHC.

In organizations having a low level of CHC with respect to the needs of the task at hand, most organizational members carry several incorrect or misaligned beliefs. The org. code will tend to imbibe a significant extent of incorrect knowledge as a consequence. When a majority of the organizational members learn slowly (i.e. under high exploration), the heterogeneity of their knowledge will decrease by acquisition of such incorrect knowledge from the org. code. A low level of organizational knowledge will be the result. In contrast, when a majority of the organizational members learn fast (i.e. under high exploitation), there will be lesser degree of degradation of knowledge heterogeneity of *Slow Learners*. This will result in a higher level of relevant organizational knowledge. Thus, in this case, higher exploitation offers more competitive advantage than will higher exploration. This suggests the following proposition:

> **P3**: When the CHC in an organization is low, the level of organizational knowledge attained by higher exploitation is greater than the level of organizational knowledge attained by higher exploration.

SIMULATION MODEL

In this study we create a simulation model from the conceptual description given in March (1991). To demonstrate the effects of varying CHC, we relax one of March's assumptions, which is that organizational members always start with random knowledge, i.e. a moderate level of CHC. We designate such populations as *Marchian*. We designate a population having a level of CHC lower than that of a *Marchian* population as a *sub-Marchian* population. Likewise, a population having a level of CHC higher than that of a Marchian population is termed as a *supra-Marchian* population. All other parts of the conceptual model described in the text of March's 1991 publication remain unchanged in our study. The environment or reality surrounding the focal organization is unchanging for our first few experiments. In the last four sets of experiments we report (Figures 6, 7, 8 and 9), we observe outcomes under varying environmental turbulence. Further, taking advantage of the advances in computing since the late-1980s, we run each experiment for 10,000 iterations (instead of the 80 iterations used in March's simulation).

Organizational Knowledge. In the model, the external reality, **R**, for the organization is an ***M***-length bit string. Each bit is considered to be a dimension of reality and can have a value of either "+1" or "–1". The values "+1" and "–1" are code numbers for different kinds of knowledge. For a given entity, say an agent or the org. code, if there is a matching value in the corresponding bit position, the entity is deemed to possess correct knowledge in that dimension **R** is initialized with a set of random values whereby each bit is equally likely to have either a "+1" or a "−1" value. The org. code, **OC**, is another ***M***-length bit string. Each dimension of **OC** can have a value from the set {"−1", "0", "+1"}. A value of "0" signifies "no opinion". **OC** is initialized with all values set to 0 at the beginning of a simulation. **OC** gets updated by learning processes as the simulation progresses. At the end of the simulation, the number of matches between the reality string, **R**, and the org. code string, **OC**, is obtained. This number, represented as a fraction of the total number of knowledge dimensions, or as a percentage, constitutes the level of organizational knowledge attained.

Collective Human Capital in an Organization. The organization is comprised of ***N*** individuals (hereafter, agents). The knowledge of each agent is modeled as an ***M***-length bit string. Comparable to the org. code, **OC**, each dimension of an agent's knowledge string can have a value from the set {"−1", "0", "+1"}. In March's model and in our baseline test at the beginning of the simulation, the knowledge string of each agent is randomly populated with values from the set {"−1", "0", "+1"}. Each value from this set has an equal probability (one-third) of materializing in any given individual agent's knowledge string. Individual knowledge strings are generated for each of the ***N*** agents. The set of ***N*** distinct knowledge strings constitutes the CHC of the organizational members (agents). The population thus generated is referred to as a *Marchian* population. For a *Marchian* population, we designate the level of CHC in the organization as *moderate*.

A *sub-Marchian* population has a higher number of occurrences of incorrect beliefs compared to a *Marchian* population (i.e. beliefs that are misaligned to the needs of the task at hand, which is defined by reality string **R**). A dimension of the knowledge of an agent is considered incorrect (or misaligned with the competitive environment) if the value for that dimension in the agent's knowledge string is non-zero but does not match with the "+1" or "−1" value in the corresponding bit position in the reality string **R**. We denote ***D*** (expressed as a percentage) as the extent of knowledge deficiency with respect to a *Marchian* population. To create a *sub-Marchian* population having deficiency ***D***, we first create a *Marchian* population and we also generate the reality string **R**. Thereafter, we overwrite ***D***% of the knowledge-string bits of each agent with values that are different from the corresponding values in the reality string **R**. For example, when the i^{th} dimension of an agent's knowledge string is to be overwritten, if the value in the i^{th} dimension of the reality string **R** is "+1", we overwrite the value in the i^{th} dimension of an agent's knowledge string with a value "−1"; conversely, when the value in the i^{th} dimension of the reality string **R** is "−1", we overwrite the value in the i^{th} dimension of an agent's knowledge string with a value "+1". The actual bits that are overwritten are randomly chosen, separately for each agent. For a *sub-Marchian* population, we describe the level of CHC in the organization as *low*.

Additionally, we designate a population having a level of CHC higher than that of a *Marchian* population as *supra-Marchian.* A *supra-Marchian* population is generated by overwriting a certain proportion of bits of the knowledge string of agents with values *matching* those in the reality string **R**. The bits that are overwritten are randomly chosen, with specific draws for each agent. Creation of a range of populations where the CHC varies enables us to observe the organizational knowledge outcomes that ensue under various configurations of learning mechanisms. For a *supra-Marchian* population, we specify the level of CHC in the organization as *high*.

Learning Mechanisms. In each time-step of our simulation runs, two concurrent learning mechanisms are operational. The org. code, **OC**, learns from the members. And each individual member also learns from the org. code. Consequently, over time, the differences in knowledge between the org. code and the members are obliterated.

Learning by the organizational code. In any given time-step, the org. code is able to identify agents who are more knowledgeable about the reality than itself. We designate these agents as elite agents. In each time-step the org. code, **OC**, identifies a set of elite agents who have relevant information needed about the organization's competitive environment and how to succeed in it. Thus, for each knowledge-bit, the org. code computes the majority opinion of the team regarding what its value should be. For example, if, for the i^{th} bit position, more elites have a value "+1" and less have value "–1", the org. code, **OC**, will deem the elites' recommendation to be for a value "+1". Zero values in the knowledge dimension of elites are ignored. If **OC** has a value other than "+1" in the i^{th} bit position, it will update it to the value recommended by the elite agents' group with a probability that is related to the code learning parameter p_2. Specifically, the probability that the i^{th} dimension of **OC** will remain unchanged at the end of the time-step is given by $(1 - p_2)^k$ with k (> 0) being the number of agents in the elite group who differ from the code in this dimension minus the number who do not (March, 1991: 74). As an illustration, when we say that the org. code learns at a rate of 10%, we mean that the value of p_2 is 0.1 ($p_2 = 0.1$).

Learning by organizational members. In each time-step, a given agent compares the value in each bit position of its knowledge string with the corresponding value in the org. code **OC**. If the value in the i^{th} position of **OC** is non-zero and different from the value in the i^{th} position of the agent's knowledge string, the agent updates its knowledge string with the corresponding org. code value with a probability p_1. The parameter p_1 constitutes the rate of learning by agents. As an illustration, when we say that a member learns at a rate of 20%, we mean that the value of p_1 is 0.2 ($p_1 = 0.2$).

The Continuum Conception of Exploration and Exploitation. The continuum conception is motivated by considering a situation where all the agents in the organization are not learning at the same rate. For modeling purposes, following March, it is assumed that a specified fraction, F, of agents learn slowly, having p1 equaling 10% (p1 = 0.1), the remainder learn fast, having p1 equaling 90% (p1 = 0.9). By suitably varying F¸ it is possible to construct populations that have average individual learning rates (p1–AVG) between 20% and 80%. When p1–AVG is 20% (p1–AVG = 0.2), about seven-eighths of the total number of members learn slowly (p1 = 0.1) and the rest learn fast (p1 = 0.9). This is mapped to maximum exploration in the continuum conception. In this situation, maximum heterogeneity is fostered in the organization by a large section of organizational members undergoing slower socialization. When p1–AVG is 80% (p1–AVG = 0.8), about seven-eighths of the total number of members learn fast (p1 = 0.9), the rest learn slowly (p1 = 0.1). This represents maximum exploitation.

Environmental Turbulence. Following March (1991) we use a parameter, p_4, to denote the degree of environmental turbulence. In stable environments the value of p_4 is set to zero. For a moderately turbulent environment, p_4 attains a value of about 2%. For more turbulent environments, $\boldsymbol{p_4}$ is given values higher than 2%. A value of $\boldsymbol{x}$% for p_4 signifies that, in any given time-step, each bit of the reality **R** has an $\boldsymbol{x}$% chance of flipping its value (from "+1" to "–1" and vice versa). In the last two experiments that we report, we invoke a punctuated equilibrium model of environmental change. In this case, the rate of environmental change is given by a parameter $\boldsymbol{p_{4\_alt}}$. A value of $\boldsymbol{x}$% for $\boldsymbol{p_{4\_alt}}$ signifies that, in any given time-step, an environmental change event materializes with $\boldsymbol{x}$% probability. Upon materialization of a change event, dimensions of the reality **R** flip values (from "+1" to "–1" and vice versa) with a probability of 50%.

Model Parameters and Robustness Checks. Following March (1991), we use a 30-bit string (M = 30) for reality, **R**, and an organization of fifty agents ($\boldsymbol{N}$ = 50). For all but the last four experiments, simulations are run for a maximum of two hundred and fifty time-steps ($\boldsymbol{T}$ = 250) for experiments in a closed system, where environmental turbulence is set to zero. The allotted maximum duration is sufficient to let the system reach equilibrium for closed systems. We report the last four experiments that incorporate environmental turbulence for tasks of duration 20 periods ($\boldsymbol{T}$ = 20), following the convention set by March (1991).[12] For any given experiment, we report results as averages over 10,000 iterations ($\boldsymbol{I}$ = 10,000). We use a 20% value for average learning rate to signify high exploration ($\boldsymbol{p_{1-AVG}}$ = 0.2); we use an 80% value for average learning rate to signify high exploitation ($\boldsymbol{p_{1-AVG}}$ = 0.8). We use a 50% level for $\boldsymbol{p_2}$, except where mentioned otherwise ($\boldsymbol{p_2}$ = 0.5). For purposes of robustness checks, we run the model with $\boldsymbol{M}$ = 25 and 35, $\boldsymbol{N}$ = 40 and 60, $\boldsymbol{p_2}$ = 0.4 and 0.6, and $\boldsymbol{p_{1-AVG}}$ values 5% on either side of high exploration ($\boldsymbol{p_{1-AVG}}$ = 0.2) and high exploitation ($\boldsymbol{p_{1-AVG}}$ = 0.8). We obtain qualitatively similar results as reported below.

RESULTS

We conducted distinct experiments, first to verify conformance of our model's mechanics to March's (1991) model—our baseline test—and then to take a deeper dive into the mechanisms underlying the continuum conception of exploration and exploitation. Towards the latter objective, we relax one of March's assumptions and allow variation in CHC embodied in the knowledge of organizational members (agents). March's assumption of random knowledge in agents at the beginning of his simulation amounted to conferring a moderate level of CHC. As stated earlier, we designate such a population as a *Marchian population*. Our baseline test with *Marchian* populations obtains qualitative conformance with March's results.[13]

Moreover, under conditions of low CHC, i.e. when organizational members initiate with a level of CHC lower than that in a *Marchian* population, a different set of outcomes transpire. These results serve to inform given conditions under which higher exploitation, rather than higher exploration is beneficial to organizations.

We conduct additional analyses to trace the exact mechanism by means of which a regime of desirability of high exploitation transitions to a regime of desirability of high exploration. We discuss our findings with respect to a set of business organizations classified according to their focus on exploration vs. exploitation, and varying in CHC.

12. Note that, under longer exposure to environmental turbulence, no particular exploration–exploitation strategy yields outcomes markedly different from that obtainable by random action.

13. We built our model according to the conceptual narration in March's published (1991) paper. Subsequently, we had access to March's code and discovered that there are some additional coding steps that are not described in the published article. Our analyses reveal that all the undocumented features should be dropped, being at a variance to the theoretical assumptions of March's paper. The dropping of non-conforming assumptions accounts for minor differences in the values obtained in our Figures 2 and 3 with respect to March's values.

Baseline Test. Figure 2 introduces the overall model mechanics. It shows the level of organizational knowledge attained (average equilibrium knowledge, on the vertical axis) when the rate of socialization of members ($\boldsymbol{p_1}$) varies (socialization rate, on the horizontal axis), for distinctive rates of code learning ($\boldsymbol{p_2}$), given that the CHC of the organizational members[14] corresponds to that present in a *Marchian* population. We observe that superior outcomes are obtained when agents learn slowly, i.e. $\boldsymbol{p_1}$ is low, and particularly so if the org. code's learning rate is fast, i.e. $\boldsymbol{p_2}$ is high. This stems from the fact that the org. code can learn only as long as its knowledge is sufficiently different from that of the members. If members learn too fast (say $\boldsymbol{p_1}$ = 0.9, on the horizontal axis, to the right), this difference is rapidly obliterated, giving less time for the org. code to learn, resulting in a lower organizational knowledge level, compared to when members learn slower (say $\boldsymbol{p_1}$ = 0.1, on the horizontal axis, to the left).

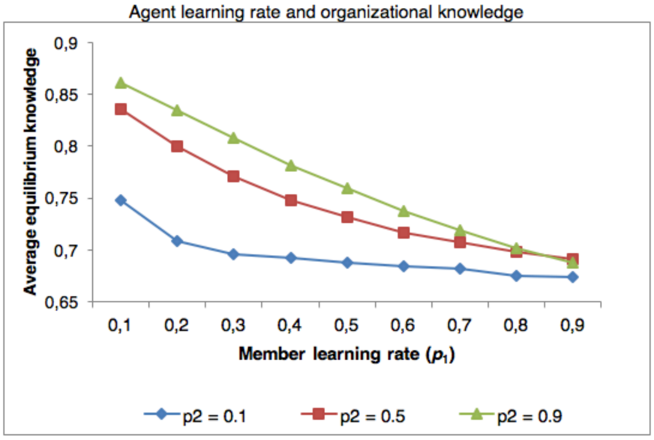

Notes. Triangular markers represent high rate of code learning ($\boldsymbol{p_2}$ = 0.9), rectangular markers represent medium rate of code learning ($\boldsymbol{p_2}$ = 0.5) and rhombus-shaped markers represent low rate of code learning ($\boldsymbol{p_2}$ = 0.1). Other parameters are $\boldsymbol{M}$ =30, $\boldsymbol{N}$ = 50, $\boldsymbol{p_4}$ =0, Iterations = 10,000, *Marchian* population.

Figure 2. Effect of learning rates (p_1, p_2) on equilibrium knowledge, *Marchian* populations

Figure 3 introduces the continuum conception. On the horizontal axis we vary average socialization rate ($\boldsymbol{p_{1-AVG}}$). Recall that a low value of ($\boldsymbol{p_{1-AVG}}$) is constructed by having a higher proportion of members learn slowly ($\boldsymbol{p_1}$ = 0.1), while the remainder of the population learns fast ($\boldsymbol{p_1}$ = 0.9) and vice versa. The characteristic for heterogeneous learning is referred to as "heterogeneous" in the graph. The characteristic labeled "heterogeneous" indicates a situation where some members of an organization learn at a slow rate ($\boldsymbol{p_1}$ = 0.1) and others learn at a fast rate ($\boldsymbol{p_1}$ = 0.9). The continuum conception is motivated by having the rate of heterogeneous learning vary from $\boldsymbol{p_{1-AVG}}$ = 0.2, signifying high exploration, to $\boldsymbol{p_{1-AVG}}$ = 0.8, signifying high exploitation. Further, in Figure 3, the characteristic labeled "homogeneous" corresponds to a situation where all members of an organization learn at a uniform rate. On the vertical axis, as before, we show an average amount of knowledge in the org. code, upon reaching equilibrium. Figure 3 shows that, for comparable average rates of learning, higher organizational knowledge is attained from heterogeneous learning. Further, organizational knowledge decreases as average learning rate ($\boldsymbol{p_{1-AVG}}$) increases from low values (20%, $\boldsymbol{p_{1-AVG}}$ = 0.2) to high values (80%, $\boldsymbol{p_{1-AVG}}$ = 0.8). The model's results suggest that, when there is a

14. The rate of socialization of members ($\boldsymbol{p_1}$) is equivalent to member learning rate. March (1991: 77) describes $\boldsymbol{p_1}$ as the socialization learning parameter.

moderate level of CHC in the organization (i.e. for *Marchian* populations), deploying a lower average learning rate, i.e. higher exploration, accomplishes higher organizational knowledge compared to higher exploitation (manifested by deploying a higher average learning rate). These results support proposition **P1**.

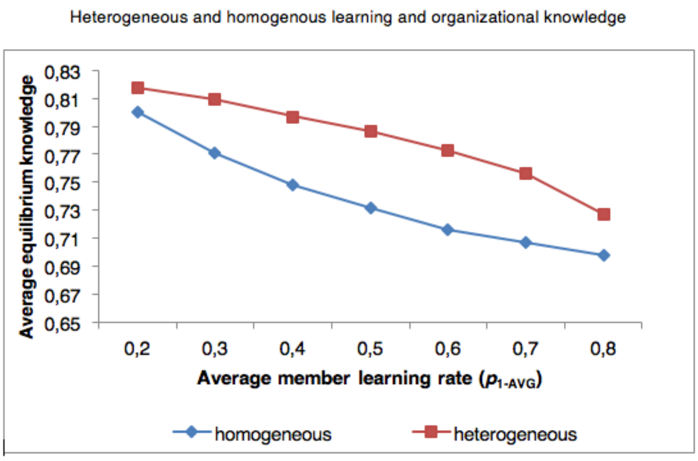

Notes. Rhombus-shaped markers represent all members learning at a homogeneous rate, rectangular markers represent a situation where some members learn at a low rate (p_1 = 0.1), others learn at a high rate (p_1 = 0.9). This characteristic, labeled heterogeneous, represents the continuum conception. Knowledge outcomes to the left are those obtained under high exploration. Those to the right are for high exploitation. The values on the *x*-axis constitute the average learning rate across members. Other parameters are M =30, N = 50, p_2 = 0.5, p_4 =0, Iterations = 10,000, *Marchian* population.

Figure 3. Effect of heterogeneous learning rates (p_1 = 0.1, 0.9) on equilibrium knowledge, *Marchian* populations

Analysis of Organizational Knowledge Outcomes under Varying Collective Human Capital. In Figure 4 we observe organizational knowledge outcomes that materialize when an organization has lower CHC compared to *Marchian* populations. We model the CHC in the organization to be somewhat misaligned to the needs of the task at hand by having 15% of randomly chosen knowledge dimensions of each member of a *Marchian* population overwritten with values opposite that of reality, **R**, at the beginning of a simulation run. As discussed earlier, when a value different from that in a dimension of reality is required to overwrite a belief dimension of an organizational member, we use “+1” if the reality string dimension has value “–1”; otherwise we use “–1”. Other parameters are identical to those in Figure 3. We observe that overall organizational knowledge levels accomplished are much lower (below 40%) than what is accomplished with moderate or *Marchian* CHC (over 70%, as shown in Figure 3). A comparison of values attained by the characteristic *heterogeneous* in Figures 3 and 4—representing the continuum conception—shows that a lower level of organizational knowledge is attained in the latter case, where CHC in the organization is low. This result supports proposition **P2**.

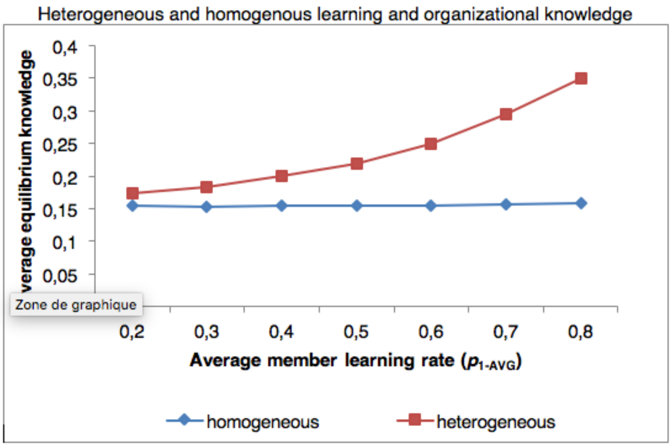

Notes. Rhombus-shaped markers represent all members learning at a homogeneous rate, rectangular markers represent a situation where some members learn at a low rate (p_1 = 0.1), others learn at a high rate (p_1 = 0.9). The values on the *x*-axis constitute the average learning rate across members. Other parameters are M =30, N = 50, p_2 = 0.5, p_4 =0, Iterations = 10,000, *sub-Marchian* population with 15% deficient human capital compared to *Marchian* populations.

Figure 4. Effect of heterogeneous learning rates (p_1 = 0.1, 0.9) on equilibrium knowledge, *sub-Marchian* populations

Furthermore, in a sharp reversal of the outcome corresponding to the continuum conception shown in Figure 3, in Figure 4 we observe that a population espousing heterogeneous learning accomplishes superior outcomes when there is a higher proportion of *Fast Learners*. Thus, when organizations initially have a low level of CHC, higher average learning rates obtain superior organizational knowledge compared to that obtained with lower average learning rates. In this situation, higher exploitation is preferable over higher exploration. Therefore, proposition **P3** is supported.
In Figure 5 we present what would be the general case underlying exploration and exploitation in the continuum conception. On the horizontal axis we vary human capital in terms of per cent offset from the level of knowledge of a *Marchian* population (CHC offset from *Marchian* population). On the vertical axis we show the level of organizational knowledge attained, at equilibrium (average equilibrium knowledge). The characteristics shown correspond to the two ends of the exploration–exploitation continuum. The characteristic with p_{1-AVG} = 0.2 indicates high exploration. The characteristic with p_{1-AVG} = 0.8 indicates high exploitation. The results in March (1991: 77, Figure 2) and our Figure 3 highlight the right-hand half of the graphs—i.e. that the high exploration characteristic (p_{1-AVG} = 0.2) results in higher organizational knowledge. This signifies that, when organizations start with moderate or high CHC, higher exploration leads to higher organizational knowledge (proposition **P1**). We show the other half of the story in the left section of the graphs—which is that higher exploitation is useful when an organization has a low level of CHC with respect to the needs of the task at hand (proposition **P3**). For a model-based explanation of the reversal of March's result upon lowering of CHC in the organization, please see Appendix A. We provide a brief summary in the next two paragraphs.

Proposition **P1** states that, for moderate CHC the outcomes of higher exploration are superior to the outcomes from higher exploitation. We observe that early in the simulation (at end of period two)—before any member learning has taken place—the org. code attains a high (40%) extent of correct knowledge, raising the bar for members who join as elites to advise the org. code. Thus, members with poor knowledge (i.e. a high proportion of incorrect knowledge) get barred from advising the org. code. In this case, the later the convergence between knowledge of all the members and the code, the higher is the level of organizational knowledge, since more of the heterogeneity of *Slow Learners* gets utilized. Hence exploration obtains outcomes superior to outcomes from exploitation.

Proposition **P3** states that for low CHC the outcomes of higher exploitation are superior to the outcomes from higher exploration. In this case, early in the simulation (at end of period two)—before any member learning has taken place—the org. code attains a high (84%) extent of incorrect knowledge. At the same time, a very low extent of correct knowledge (7%) in the org. code allows members with highly faulty knowledge to get into the team of elites who advise the code. The longer the exposure of the members to the highly error-ridden org. code, the greater is the knowledge lost from all members and particularly so, from the slow-learning members. For a high rate of exploitation, the duration of exposure is lower. Hence knowledge of the slow-learning members degrades less. This translates to superior outcomes for high exploitation.

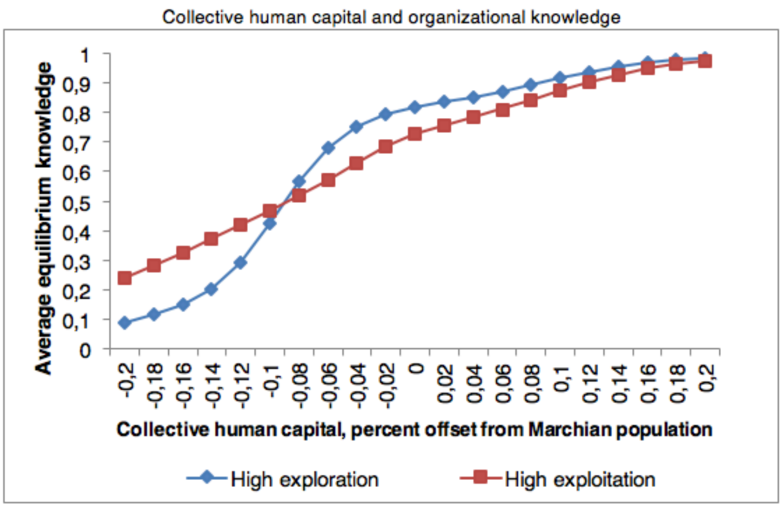

Notes. Rhombus-shaped markers represent members learning at an average rate of 20% ($p_{1\text{-AVG}}$ = 0.20), rectangular markers represent members learning at an average rate of 80% ($p_{1\text{-AVG}}$ = 0.80). In both cases, some members learn at a low rate (p_1 = 0.1), others learn at a high rate (p_1 = 0.9). The values on the *x*-axis constitute offset from the level of collective human capital of *Marchian* populations. Negative values signify misalignment with needs of the task at hand. Positive values signify alignment. Other parameters are M =30, N = 50, p_2 = 0.5, p_4 =0, Iterations = 10,000.

Figure 5. Effect of variation in collective human capital on equilibrium knowledge

The results shown in Figure 5 illustrate a form of phase transition upon reaching critical mass. On one side of the critical point, the winning strategy concerns high exploitation. On the other side, the winning strategy concerns high exploration. In Figure 5, the critical point is situated where the two characteristics (plot-lines) cross. This happens for a level of CHC at an offset of approximately 10% below the knowledge of *Marchian* populations. The exclusion of cases where the CHC is lower than the human capital embodied in *Marchian* populations leads to a conclusion that more exploration is always better, calling into question the notion of balance between exploration and exploitation. Dierickx & Cool (1989) note that joint consideration of stocks and flows is necessary to develop adequate theory regarding phenomena. By associating the stocks of CHC and organizational knowledge in an organization with the flow elements comprising of organizational learning rate and member learning rate, we advance March's (1991) findings to the next logical step in theory development.

Organizational Knowledge Outcomes under Varying Environmental Turbulence. In Figure 6 we present organizational knowledge outcomes attained by *sub-Marchian* populations for a range of environmental turbulence. In Figure 7 we present the corresponding information for *Marchian* populations.

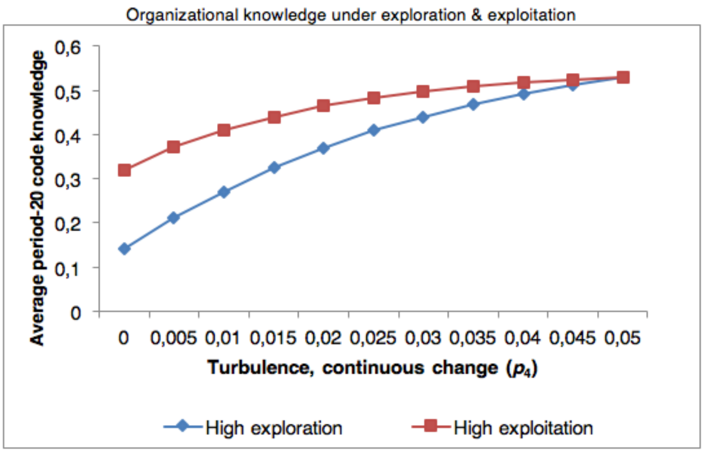

Notes. Rhombus-shaped markers represent members learning at an average rate of 20% ($p_{1\text{-AVG}}$ = 0.20), rectangular markers represent members learning at an average rate of 80% ($p_{1\text{-AVG}}$ = 0.80). In both cases, some members learn at a low rate (p_1 = 0.1), others learn at a high rate (p_1 = 0.9). Other parameters are M =30, N = 50, T = 20, p_2 = 0.5, Iterations = 10,000, *sub-Marchian* population.

Figure 6. Effect of variation in environmental turbulence on organizational knowledge, s*ub-Marchian* populations, continuous change environment

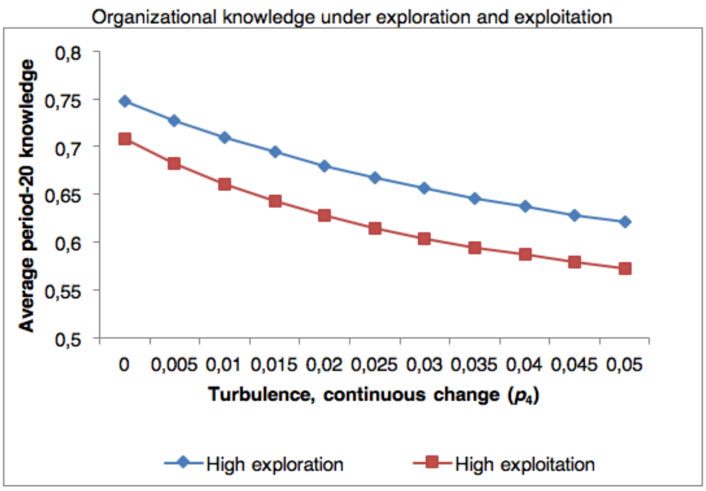

Notes. Rhombus-shaped markers represent members learning at an average rate of 20% ($p_{1\text{-AVG}}$ = 0.20), rectangular markers represent members learning at an average rate of 80% ($p_{1\text{-AVG}}$ = 0.80). In both cases, some members learn at a low rate (p_1 = 0.1), others learn at a high rate (p_1 = 0.9). Other parameters are M =30, N = 50, T =20, p_2 = 0.5, Iterations = 10,000, *Marchian* population.

Figure 7. Effect of variation in environmental turbulence on organizational knowledge, *Marchian* populations, continuous change environment

We observe that the results noticed for a stable environment by and large continue to hold good for low to moderate turbulence (0% to 2% turbulence as per the convention we adopted). Carrying out higher exploitation (focusing on execution) continues to be a better strategy than engaging in higher exploration (experimentation), for *sub-Marchian* populations representing organizations low in CHC. The differences narrow and disappear for higher turbulence. For *Marchian* populations—organizations that are not deficient in CHC—higher exploration (experimentation) continues to obtain better results, compared to exploiting more (being focused on execution and refinement). The differences persist for the entire range of environmental turbulence examined.

The result presented above is different from the findings of Posen and Levinthal (2012). Posen and Levinthal's (2012) single-actor, multi-armed, bandit model suggests that a 50:50 mix of exploration and exploitation is optimal in a stable environment (p. 592, Figure 1). In a changing environment, higher exploitation is favored (p. 594, Figure 3). Our results are different because our formal model allows utilization of heterogeneous knowledge dispersed in the heads of organizational members (CHC), something a single-actor model cannot do. We further observe that Posen and Levinthal (2012) invoke a *punctuated equilibrium model of environmental change*, unlike March (1991) where turbulence is characterized by *continuous environmental change*. In the punctuated equilibrium model of environmental change, upon materialization of a change event—say with probability $\boldsymbol{p_{4\_alt}}$—dimensions of reality flip values (from "+1" to "–1" and from "–1" to "+1") with 50% probability. In Figures 8 and 9 we present results analogous to those in Figures 6 and 7 respectively, but using a punctuated equilibrium model of environmental change instead of a continuous change model. We observe that the results are qualitatively similar to that in Figures 6 and 7. The gap between the characteristics for high exploration and high exploitation is narrower in Figures 8 and 9. Again, focusing on execution (exploitation) comes out as a desirable strategy for organizations with *sub-Marchian* population, and focusing on experimentation (exploration) is observed to work better for organizations with *Marchian* populations.

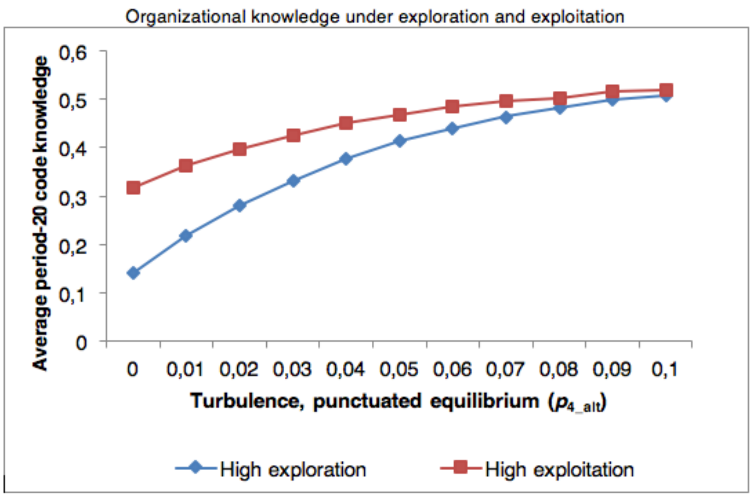

Notes. Rhombus-shaped markers represent members learning at an average rate of 20% ($\boldsymbol{p_{1\text{-AVG}}}$ = 0.20), rectangular markers represent members learning at an average rate of 80% ($\boldsymbol{p_{1\text{-AVG}}}$ = 0.80). In both cases, some members learn at a low rate ($\boldsymbol{p_1}$ = 0.1), others learn at a high rate ($\boldsymbol{p_1}$ = 0.9). Other parameters are $\boldsymbol{M}$ =30, $\boldsymbol{N}$ = 50, $\boldsymbol{T}$ = 20, $\boldsymbol{p_2}$ = 0.5, Iterations = 10,000, *sub-Marchian* population.

Figure 8. Effect of variation in environmental turbulence on organizational knowledge, s*ub-Marchian* populations, punctuated equilibrium model of environmental change

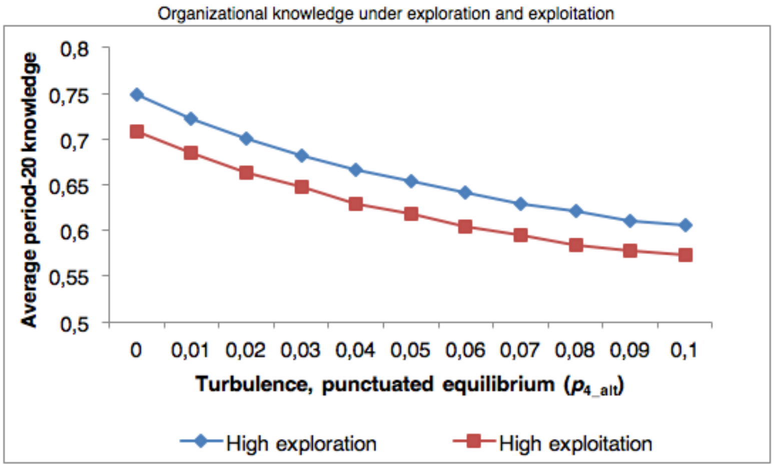

Notes. Rhombus-shaped markers represent members learning at an average rate of 20% ($p_{1\text{-AVG}}$ = 0.20), rectangular markers represent members learning at an average rate of 80% ($p_{1\text{-AVG}}$ = 0.80). In both cases, some members learn at a low rate (p_1 = 0.1), others learn at a high rate (p_1 = 0.9). Other parameters are M =30, N = 50, T =20, p_2 = 0.5, Iterations = 10,000, *Marchian* population.

Figure 9. Effect of variation in environmental turbulence on organizational knowledge, *Marchian* populations, punctuated equilibrium model of environmental change

Results Summary. We present a summary of our results in Table 2. The lowest level of organizational knowledge is attained when an organization with low CHC pursues a high degree of exploration (Quadrant I, labeled as Q1 in the Table). For example, from Figure 4 we find the value to be around **17%**, as indicated by the value attained by the curve *heterogeneous* at the left-most point, where p_{1-AVG} = 0.2. Somewhat higher organizational knowledge is attainable when an organization with low CHC exploits at a high rate (Quadrant II, labeled as Q2 in the table). For example, from Figure 4, we find the value to be around **35%**, as may be observed from the value shown by the curve *heterogeneous* at the right-most point, where p_{1-AVG} = 0.8. If an organization possesses moderate or high CHC, it stands to attain even higher organizational knowledge by high exploitation (Quadrant III, labeled as Q3 in the Table). For example, from Figure 3, we find the value to be around **70%**, as indicated by the value attained by the curve *heterogeneous* at the right-most point p_{1-AVG} = 0.8. The highest level of organizational knowledge is attained, however, when an organization with moderate or high human capital explores at a high rate (Quadrant IV, labeled as Q4 in the table). This may be observed from Figure 3, in the curve *heterogeneous* at the left-most point where p_{1-AVG} = 0.2, which attains over **80%** organizational knowledge. The relative orders of magnitude are noteworthy: deficient-CHC firms can obtain 2X the organizational knowledge by pursuing high exploitation instead of high exploration. Further, firms not deficient in CHC obtain upwards of 2X the maximum attainable by deficient-CHC firms.

	Organizational Knowledge attained under **High Exploration**	**Organizational Knowledge** attained under **High Exploitation**
Low collective human capital	Very Low (Q1)	Low (Q2)
Moderate or high collective human capital	Very High (Q4)	High (Q3)

Table 2. Relative levels of organizational knowledge attainable by high exploration and high exploitation, for organizations with low moderate and high collective human capital

Interpretation of the Results. Business organizations or companies in Q1 and Q2 are those that have (a) fairly low levels of CHC and (b) negligible access to external resources. Exploration is a luxury that companies in this situation cannot afford. They survive by selling low-cost, lean-feature (simple) products. A successful firm operating in a bottom-of-pyramid (BoP) market constitutes a good example of this kind of firm. If the relative level of organizational knowledge attainable is an indicator of probability of success in a given market, companies in Q1 are unlikely to survive for long. Companies in Q2, on the other hand, may find vacant niches and under-served markets and then develop into significant businesses over time.

The existence of companies in Q2 gives rise to an interesting phenomenon frequently encountered in developing economies. A newcomer to an industry with fairly low initial resources and human capital (a) creates a feature-lean product at a price much lower than the prices of its competitors and (b) rapidly attains an efficient scale of operations. Low CHC is indicated by the fact that such a firm has limited means. A lower level of organizational knowledge (compared to mainstream businesses) is manifested in the feature-lean products and simple distribution strategies the firm comes up with. Over time, this company goes on to emerge as a formidable competitor to the industry incumbents. For a good example, please see ICMR (2001), The *Nirma* Story. It describes how a detergent brand of humble origins—initially sold by the maker of *Nirma* from his own bicycle—took on Unilever's powerful subsidiary in India. Lower CHC limits such businesses to only a handful of features and functionality. In the process though, a new segment opens up, at a price point far lower than what the incumbents can anticipate. In some cases, over time, a viable business emerges—one with cash positions sufficient to make forays into more mainstream products and services.

The organizations in quadrant Q3 are ones that do not lack human capital. However, their primary mode of operating is by cutting costs (high exploitation). Developed-economy (for example US-based) firms that have outsourced manufacturing to developing country locations (for example China) fall in this category. The organizations in Q4 also do not lack human capital. However, they focus more on innovation and less on cutting costs. A firm like Intel, which leads in innovative design of new products and endeavors to make continuous improvements in semiconductor manufacturing, falls in this category. As long as their innovations are up to date and useful, the Q4 organizations may be expected to do better than Q3 organizations (note, however, that as we write this paper, Intel is suffering from a lack of new chip-related inventions).

DISCUSSION

Our research was motivated by a feeling of puzzlement that, even though the necessity of balancing exploration and exploitation is widely accepted, March's formal model—considering exploitation and exploration as two ends of a continuum in a closed system—appears to suggest that higher exploration is always preferable over higher exploitation. We unearth an additional principle by considering the knowledge stocks (CHC and organizational knowledge) jointly with the flows (comprising member and org. code learning rates). Thus, in organizations deficient in CHC, high exploitation is preferable over high exploration. In the process, we add a boundary condition to March's finding that high exploration (experimentation) is preferable over high exploitation (execution and refinement focus) only when the CHC in an organization is moderate or high and *not* when the CHC in the organization is low. These two principles apply for an organization with negligible access to heterogeneous knowledge from outside and operating under low to moderate environmental turbulence.

Our study, per se, does not identify a ratio of exploration–exploitation that is ideal. Therefore, balancing exploration and exploitation—in the continuum sense—may mean that an organization allocates the minimal acceptable resources to one activity and concentrates the rest of its resources on the other activity (Levinthal & March, 1993). The role of managerial judgment is crucial in determining the level that constitutes the minimum for the less-preferred activity.

Implications for Research. There are important implications of the fact that the formal model for the continuum conception fails to demonstrate superiority of balancing exploration and exploitation in an organization. Our computational model clarifies that the formal model actually suggests that lack of balance is not an impediment to attaining superior organizational knowledge. *The level of CHC endowment of an organization determines whether it is rewarding to focus on exploration or exploitation*. This principle suggests that there is scope for managerial intentionality (Hutzschenreuter. Pedersen & Volberda, 2007; Chanda and Ray, 2015) towards orienting an organization either to exploratory innovation or to exploitative innovation (Jansen, Van den Bosch and Volberda, 2006; Chanda & Ray, 2015) under an assumption that that the level of CHC can be chosen by organizations—for example through acquisitions and divestments. Empirical evidence exists that companies can be successful with "unbalanced" orientation: for example, Chen and Katila (2008) cite IBM and Dell as leaning towards high exploitation, and Amazon and Apple as leaning towards high exploration. Moreover, upon orienting their organization towards exploratory or exploitative innovation (by configuring an appropriate level of CHC), managers have to commit to carrying out a certain minimum level of the non-preferred activity. There is a need for further research on what that minimum level should be for organizations operating in particular contexts. The new "balance discussion" needs to focus on finding the determinants of the minimum level of the non-preferred activity for organizations into exploratory innovations as well as for organizations into exploitative innovation.

We provide a different answer to the question as to what proportion of exploration and exploitation is desirable in an organization. Posen and Levinthal (2012) use a single-actor model to show that a 50:50 distribution is optimal, in a stable environment (p. 592, Figure 1); for dynamic environment, the proportion changes, eventually favoring higher exploitation under higher turbulence (p. 594, Figure 3). By investigating the implications of the continuum conception of exploration and exploitation in the formal model proposed by March (1991), our study finds little justification for altering a successful exploitation–exploration mix—for companies deficient in CHC as well as companies not deficient in CHC—in a changing environment.

A single-actor model is ill-suited to suggest outcomes of efforts by a group of organizational members. In the Posen and Levinthal (2012) study, a single agent updates estimates of probabilities of obtaining payoff by operating the levers of a multi-armed bandit model successively over time. Learning is by one agent, starting from zero knowledge. In March (1991) and our model, the org. code also initiates with zero knowledge. Thereafter, the genetic algorithm mechanism allows rounding up of heterogeneous knowledge dispersed in the heads of organizational member agents—CHC—as validated knowledge in the org. code. CHC is operationalized as a property of a group, where the whole is not the sum of the parts. Rather, CHC is fashioned from the aligned, misaligned and neutral beliefs regarding organizational matters, embedded in the latent knowledge of the organizational population. Thus, we substantiate our distinctive answer from a study of the emergence of a macro-level pattern of organizational knowledge from the micro-level learning behavior of organizational members.

Implications for Theory. Our study brings to the fore the important role of CHC in an organization for determining whether it should focus on exploitation at the expense of focusing on exploration or vice versa. Further research is necessary to investigate how a company senses the status of its own human capital and determines the uses that the latter may be deployed for.

Second, by linking the effectiveness of managerial choice—with regard to allocation of resources towards exploration and exploitation—with the extent of CHC in the organization, we offer an explanation to the puzzling phenomenon that higher exploration is deemed as desirable (as from March's existing theory), but is not undertaken in some companies. These companies seem to be happy to exploit on a large scale. A conjunction of (a) minimal accessing of outside resources and (b) a low level of CHC motivates the wisdom behind choosing high exploitation over exploration.

Third, our study highlights that successful strategies—focusing on exploration vs. focusing on exploitation—are different, depending on whether an organization possesses a critical mass of CHC or otherwise, respectively. Further research is necessary to find factors that determine the extent of CHC that would qualify as critical mass in stylized contexts. For example, if one company A acquires another company B, does the CHC automatically increase, for the merged entity? Moreover, post-acquisition, if many organizational members leave, what factors determine the danger-mark of falling below the critical mass of CHC? Alternatively, if a merger or a divestment is in response to increasing commoditization of an industry's products (or services), what determines the extent of CHC that a company may safely shed in order to transition to a regime of high exploitation? Conversely, what determines the measures for increasing and nurturing CHC, if a company, upon amassing large cash reserves through success in serving BoP markets with feature-lean products wishes to enter the mainstream? Is entry to the mainstream with a high exploitation focus the only option? Or, by suitable nurturing of CHC, can a company leapfrog to offering differentiated products in the mainstream directly?

Fourth, our study suggests a mechanism by which organizations low in CHC may win over other similar organizations, i.e. by purposefully choosing to exploit at a high rate. Such action opens up the opportunity to serve new markets for low-cost, feature-lean products and services. Our work could be applied to understanding how a company may obtain success serving BoP markets. Though much has been written already (ICMR, 2001; Prahalad, 2004; Prahalad & Hart, 2002), a theoretical treatment with the exploration–exploitation paradigm has hitherto escaped the attention of scholars. Our study breaks ground in this direction.

A company lacking in CHC endowment deploys technology that is low in sophistication to manufacture feature-lean products. Most often, such products are welcome only in the BoP markets. Moreover, the company follows simple sourcing and distribution strategies, thereby making do with a relatively lower level of organizational knowledge, compared to mainstream incumbents. For example, the originator of the Nirma detergent engaged poor women working at home to pack his produce. He himself distributed the product—initially only to friends and relatives en route to his office—riding on his bicycle.

For products targeted to the BoP customers, the price is set well below offerings from more renowned firms. Yet, the minimalist functionality that is delivered reflects deep understanding of the needs and aspirations of the target customers. The rate of exploitation picks up as the entrepreneurial company scales up its cost-effective ways to serve needs left unmet by mainstream businesses. Moreover, a certain degree of exploration is involved in divining newer ways to reach customers, forging supply-chain relationships with members of the BoP community, and so forth. However, these companies cannot afford costly experimentation. Unlike better-endowed mainstream businesses, they have limited ability to bounce back from failure in experimentation. A high rate of exploration is likely to be fatal for these firms.

Our study's most important message is that all companies should not strive to do more exploration all the time. Companies low in CHC are better off focusing on high exploitation. From a consulting perspective, going forward, we refrain from suggesting ways and means for increasing exploration for this class of companies. Further, when companies shed CHC—say by divestiture or spin-off—the entity that remains may have to critically evaluate its CHC, to determine whether a higher exploitation focus is more suitable after the organizational change. Where a change in focus is applicable, additional tasks get added to the organizational change management efforts.

Limitations of our Study. The validity of the propositions from the model is crucially dependent on the proper functioning of the mechanism that the org. code obtains new knowledge only from the elite members (i.e. agents who are more knowledgeable about the reality relevant to the organization than the org. code). In the event an organization functions by learning from all members by a democratic process, there is no impact of varying member learning rate on organizational knowledge.[15] In such a case, it is not possible to distinguish exploration from exploitation by means of the average rate of learning by members. Further, we have limited the scope of the study to leveraging knowledge heterogeneity internal to a company. Future research may consider implications of permitting inflow of knowledge from outside the system, perhaps by hiring new employees who bring new knowledge into an organization by means of strategic alliances, and so forth.

Moreover, the genetic algorithm computational model, by design, ignores the effect of learning between members on increasing organizational knowledge. The research by Miller, et al. (2006) and Miller and Martignoni (2016) holds out the promise of exciting new avenues when the model is extended to allow interpersonal learning, and particularly so when tacit knowledge is involved.[16]

Future Research Opportunities. From our study, the level of CHC in the organization gets identified as a crucial parameter that determines which of exploration or exploitation ought to be preferred, given an organization's resource capabilities and environmental constraints. Thus, there is need to develop understanding of factors that influence the location of the critical point at which the switchover happens. As a small step in this direction in Appendix B we show how the critical level of CHC varies with an increase in the values of two model parameters—the number of members, signifying size of an organization, and the number of

15. The results of separate experiments conducted to verify this point are available from the authors upon request.

16. Miller, et al. (2006) allow members to learn codifiable knowledge from the organizational code. However, for non-codifiable knowledge (or tacit knowledge) they consider interpersonal exchange a necessary requirement.

dimensions of reality, signifying the breadth of context. In both cases we observe that the critical level of CHC moves up closer to the value in *Marchian* populations.

Moreover, in the new "balance" discussion that we identify, organizations carry out a minimum level of one activity from exploration and exploitation and concentrate resources on the other activity. We note that a company will surely fail if it carries out less than a minimum level of a non-preferred activity. At the same time, there is scope for gains when a company realizes it is doing more than the minimum necessary and reallocates resources as a result. Future research may find it productive to investigate how managers form judgments regarding the minimal level of the less-preferred activity and what environmental and organizational conditions cause the desirable level to change over time.

We also note that, in March's model, changes toward more or less exploration and/or exploitation levels are assumed to occur in a costless fashion. In real-life cases, there may be considerable costs to lower one activity and/or increase emphasis on the other. Incorporation of related constructs and processes holds out the promise of new theory.

Moreover, in multi-business firms, individual units comprise distinct aggregates of CHC, each directed to its specific goals. Future research may fruitfully inquire what patterns emerge when such collectives interact in the organizational ecosystem. Alternatively, even a single-business firm is embedded in a network of relationships with suppliers, customers, alliance partners and other entities. While the status of its own CHC may bias a firm towards maximizing outcomes by either exploration or exploitation, feedback from partners may work in unexpected ways to re-orient the action of such firms. Contingencies other than the status of the CHC of an organization that shape a company's orientation towards exploration or exploitation constitute an interesting area for further study.

CONCLUDING REMARKS

In this section we take stock of our progress in realization of March's wisdom that "ideas that transform ways of thinking about practical problems rarely come from a direct focus on those problems" (2013: 732).
First, our endeavor to update March's (1991) theory by joint consideration of stocks and flows led us to discover that the formal model for the continuum conception of exploration and exploitation does not *demand* that the two activities be balanced in order to produce superior organizational outcomes. This necessitates a new discussion on "balance" where the focus of inquiry is on determination of the minimum level of the less-preferred activity (from exploration and exploitation) to be carried out, given managerial intentionality towards either exploratory or exploitative innovation. Second, our research enables us to provide a different answer to a question recently addressed by noted scholars in the field—what mix of resource allocation for exploration and exploitation is desirable for stable and dynamic environments? Third, we suggest a theoretical explanation of an under-theorized phenomenon, that firms having limited means may rise to prominence, by serving BoP markets with bare-bones products, avoiding experimentation (these firms have limited ability to continue existing after failures in experimentation) and instead focusing on excelling in a limited set of activities (as evidenced by the attainable organizational knowledge range of 15-35%), when regular firms (that are not low in CHC) operate with organizational knowledge twice that level.

Last but not the least, in a bid to explain why successful strategies are different on either side of a *critical* level of CHC, we develop a thorough explanation of a phenomenon that is akin to phase transition—as in the terminology of complexity theorists—given that the maximum organizational knowledge on the deficient-CHC side is more than twice that reached on the moderate-CHC side. The *Cambridge Dictionary* defines *critical mass* as "the size that something needs to reach before a particular change, event or development can happen".[17] The explanation provided in Appendix A concerning how a small change in CHC leads to a significant change in the level of organizational knowledge attainable is probably the first that goes to micro-level dynamics to explain a phenomenon similar to phase transition upon reaching (or falling below) critical mass. In the world of finance, the term critical mass is deployed to signify "the point at which a growing company becomes self-sustaining, and no longer needs additional investment to remain economically viable".[18] In our research we show *how* a slight change in a stock variable (level of CHC in the organization) leads to a large change in attainable outcome (from less than 35% to over 70%) as well as reversal of the desirable strategy (the winning formula shifts from high exploitation to high exploration). We draw the attention of scholars and practitioners to the mechanism by which the radical impact of relatively small change to concentration to certain stocks comes to be. Thereby, we pave the way for further research inquiry into phenomena involving critical mass effects.

17. In Physics, the term "critical mass" signifies "the smallest amount of matter that is needed to produce a nuclear chain reaction" (*Cambridge Dictionary*, https://dictionary.cambridge.org/dictionary/english/critical-mass. Accessed: March 29, 2018).

18. Source: https://www.investopedia.com/terms/c/critical-mass.asp. Accessed: March 29, 2018.

APPENDIX A

Model-based explanation of the result suggesting strategies emphasizing exploitation are advantaged compared to strategies emphasizing exploration when level of collective human capital (CHC) is low, and disadvantaged otherwise

In order to develop an appreciation as to why, in the formal model for the continuum conception of exploration and exploitation, (a) greater exploration is favorable when CHC is moderate (say at the level represented by *Marchian* populations) or higher, and (b) greater exploitation is favorable when CHC is low (say at the level represented by *sub-Marchian* populations) or lower, we need to study how knowledge held by the org. code and organizational members' beliefs change over time.

WE ADOPT THE CONVENTIONS THAT:

(a) When we refer to the contents of the org. code, we refer to knowledge being "correct" if the value in a given position matches that in the corresponding position of the reality (**R**) or environment. Further, knowledge bits with "0" values do not contribute in the count of correct or incorrect knowledge. The other dimensions that host non-matching bits between the reality and org. code are referred to as *incorrect knowledge*.

(b) When we refer to the contents of the members' belief strings we cite a belief as correct if it matches with the value in the corresponding dimension of the reality string; otherwise, we cite that as *misaligned* or *wrong beliefs*. Belief bits with "0" values do not contribute to the count of correct or incorrect beliefs.

Given the rather involved nature of the explanation, we adopt the following strategy to present our explanations. First, we make some analytical predictions—by means of logical arguments—regarding what values of member and organizational knowledge we expect to see in the first two time-steps. Next, we check the graphical output having org. code knowledge as the dependent variable to see whether our analytical predictions fructified. Upon obtaining confirmation, we make a second set of predictions regarding what we expect to see in time-steps beyond the second time-step. The second set of predictions is regarding how knowledge of the members of the organization take shape.[19] Throughout, we highlight *statements that are analytically derived* by underlining them. We also assign numbering in Roman numerals to identify the analytical statements. This recourse helps demarcate analytically derived inferences from other text and description. When the analytically predicted outcomes show up in graphical output, we obtain confirmation about the line of reasoning and draw the attention of the reader to the confirmation by referring back to the analytical statement identified by Roman numerals (I, II … X).

ANALYTICAL PREDICTIONS ABOUT MODEL BEHAVIOR IN FIRST TWO TIME-STEPS

Recall that a *Marchian* population is constructed by randomly assigning "–1", "+1" and "0" to each of M belief dimensions of each of N members of the organization. Each value has one-third probability of

19. To reiterate, we make predictions regarding *organizational code knowledge* in the first two periods by logical arguments. We confirm these subsequently by checking the graphical output. At this point we make predictions regarding the *knowledge of organizational members (Slow and Fast Learners)* for subsequent time-steps. Thereafter we verify that the analytical predictions were indeed borne out.

materializing in any member's belief string. Also recall that the external reality or environment string is formed by randomly assigning "–1" and "+1" values, with 50% probability, to each dimension. Thus, on average, two-thirds of member bits have value "–1" and "+1". Thus, a _*Marchian* population, on average, contains one-third (33%) correct beliefs_ **[I]**. Further, recall that a *sub-Marchian* population (in Figure 4) was constructed by randomly assigning 15% of member bits values opposite to that in the reality string.[20] We expect a third of these 15% bits to have overwritten "0" values, another third to have overwritten correct knowledge bits, and the last third to have overwritten incorrect (or misaligned) beliefs of members.[21] Thus, a _*sub-Marchian* population initiates with (33 – 5) % = 28% correct beliefs_, on an average **[II]**. Following similar reasoning we assess that a *Marchian* population has one-third (33%) incorrect (or misaligned) beliefs on an average, whereas a *sub-Marchian* population has an average of (33 + 10) % = 43% incorrect (or misaligned) beliefs at the time of start of simulation runs.

We also need to appreciate a fine point regarding the cadence of updates to org. code knowledge and member belief values. In the March (1991) model, in each time-step, the org. code learns from members and members learn from the code. In the program (computer code) implementation, the organizational members learn from the image of the org. code that existed in the immediately prior time-step. Likewise, the org. code learns from the image of the member belief strings that prevailed in the immediately previous time-step. This is important. Let us see why.

Recall that, at time $t = 0$, the org. code has value "0" (no opinion) in all dimensions. This has two important implications.

First, in time-step *one* ($t = 1$), every member will qualify as a member of the elite group the org. code learns from, since all have one or more than one bit matching with reality.[22] Thus, the org. code knowledge that we observe in time-step *two* ($t = 2$) is an outcome of the org. code learning from all members for one period. Hence, given that *Marchian* and *sub-Marchian* populations initiate with different average correct beliefs (33% and 28% in our example), we expect the org. code knowledge to be different in time-step *two* ($t = 2$), between *Marchian* and *sub-Marchian* populations. We expect _higher knowledge in org. code in time-step *two* ($t = 2$) for *Marchian* populations, than for *sub-Marchian* populations_ **[III]**. Further, since *sub-Marchian* populations have members with a higher number of wrong beliefs compared to *Marchian* populations (43% vs. 33%) we expect the _org. code to have a higher proportion of incorrect knowledge for *sub-Marchian* populations_ compared to *Marchian* populations in time-step *two* ($t = 2$) **[IV]**.

A second implication of the org. code having all "0" values in the first time-step is that the organizational members learn nothing from the org. code in time-step *one*. Thus, the average knowledge of organizational members remains the same in time-steps *one* ($t = 1$) and *two* ($t = 2$). There is no effect of the rate of member learning on the member knowledge being observed in time-steps *one* ($t = 1$) and *two* ($t = 2$). The first learning episode that changes member knowledge values occurs at the end of time-step *two* ($t = 2$). The results of such learning on org. member beliefs are observed from time-step *three* ($t = 3$) onward.

With the above background, we set out to explain the reversal from desirability of high exploration to desirability of high exploitation upon

20. In the construction of a *sub-Marchian* population, if a member bit is chosen for update, a value of "+1" is given if the corresponding reality bit is "–1". Else, a value of "–1" is given. Reality can have only "+1" and "–1" values.

21. We note that, the computer code process of overwriting produced no change in value for the last third of member bits that are overwritten, since those had incorrect beliefs anyway.

22. This holds true except in the rarest of rare cases, viz. all members initiate with not a single belief dimension matching the corresponding dimension in the reality **R**, just by chance. This case can be safely ignored, given we carry out 10,000 iterations and observe knowledge values that are averages across 10,000 runs (Monte Carlo technique).

change of CHC from moderate level (as in a *Marchian* population) to low level (as in a *sub-Marchian* population). The explanation has three components. First, we look at change in org. code knowledge over time. Next, we look at change in the proportion of incorrect knowledge in the org. code over time. Particularly, we focus on the value of org. code knowledge and the proportion of incorrect knowledge bits in the org. code at the end of time-step *two* (t = 2), a point up to which nil member learning has taken place. Thereafter we look at the change in knowledge of *Slow Learner* and *Fast Learner* organization members over time.

CHANGE IN ORGANIZATIONAL CODE KNOWLEDGE OVER TIME AND PREDICTIONS REGARDING MEMBER KNOWLEDGE IN SUBSEQUENT TIME-STEPS

In Figure A1A we present org. code knowledge over time for high exploration (p_{1-AVG} = 0.2) and for high exploitation (p_{1-AVG} = 0.8) for an organization having moderate CHC (i.e. a *Marchian* population). In Figure A1B we present comparable information for an organization having low CHC (i.e. a *sub-Marchian* population). In Table T1 we present the numerical values of org. code knowledge underlying these graphs pertaining to the first six time-steps, in lines L1 … L4.

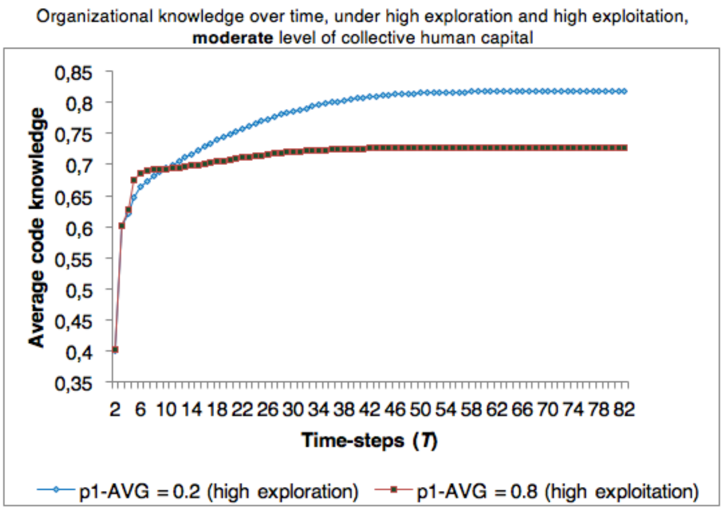

Notes. Rhombus-shaped markers represent members learning at an average rate of 20%, rectangular markers represent members learning at an average rate of 80%. In both cases, some members learn at a low rate (p_1 = 0.1), others learn at a high rate (p_1 = 0.9). Other parameters are M =30, N = 50, p_2 = 0.5, iterations = 10,000, *Marchian* population.

Figure A1A. Effect of heterogeneous learning rates ($p_{1\text{-}AVG}$) on organizational code knowledge over time, *Marchian* populations

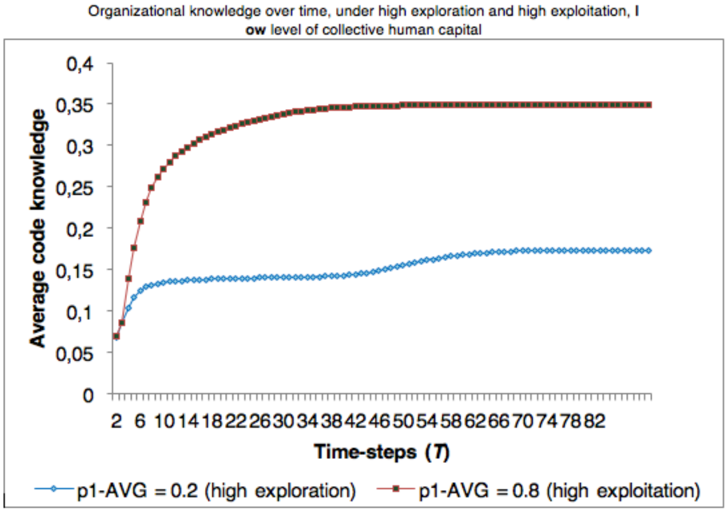

Notes. Rhombus-shaped markers represent members learning at an average rate of 20%, rectangular markers represent members learning at an average rate of 80%. In both cases, some members learn at a low rate (p_1 = 0.1), others learn at a high rate (p_1 = 0.9). Other parameters are M =30, N = 50, p_2 = 0.5, iterations = 10,000, *sub-Marchian* population with deficiency 15% from *Marchian* population

Figure A1B. Effect of heterogeneous learning rates ($p_{1\text{-AVG}}$) on organizational code knowledge over time, *sub-Marchian* populations

SI.	**Collective human capital**	**Org. code knowledge for** High Exploration ($p_{1\text{-AVG}}$ = 0.2)						
	Time-Step→	***t* = 1**	***t* = 2**	***t* = 3**		***t* = 4**	***t* = 5**	***t* = 6**
L1	Moderate (*Marchian*)	0	0.401	0.602		0.621	0.647	0.663
L2	Low (*sub-Marchian*)	0	0.069	0.086		0.104	0.117	0.125
		Org. code knowledge for High Exploitation ($p_{1\text{-AVG}}$ = 0.8)						
L3	Moderate (*Marchian*)	0	0.403	0.602		0.626	0.676	0.685
L4	Low (*sub-Marchian*)	0	0.069	0.086		0.139	0.176	0.209

Table T1: Org. code knowledge values in the first six time-steps, for high exploration and high exploitation, for moderate collective human capital and low collective human capital

From lines L1 and L3 of Table T1 we learn that, in time-step *two* (*t* = 2), the org. code attains about 40% correct knowledge, for moderate human capital (*Marchian* populations). Recall that, up until this point no member has acquired any knowledge from the org. code. Also, note that only members with more than 40% correct beliefs get to be part of the elite the org. code *subsequently* learns from.[23] This lowers the likelihood of the code obtaining incorrect knowledge subsequently. We further reason that, given that org. members initiate with 33% correct knowledge, and given that the org. code imposes a bar of 40% to constitute the elite group that the code learns from, <u>the knowledge of the organization's members goes up over time till equilibrium is reached</u> **[V]**.

23. By design, the org. code learns from a group of elites comprising of members who have more knowledge than itself.

In contrast, lines L2 and L4 of Table T1 inform that, in time-step *two* (t = 2), the org. code attains about 7% correct knowledge, when CHC is low (*sub-Marchian* population).[24] Thus, in this case, the bar for getting into the elites—from whom the org. code learns in subsequent time-steps—is much lower than for *Marchian* populations. This suggests that in *sub-Marchian* populations, the org. code will be infiltrated with a higher proportion of incorrect knowledge, compared to *Marchian* populations **[VI]**.[25] Further, note that members of a *sub-Marchian* population initiate with about 28% correct knowledge. However, they learn from an org. code with only 7% correct knowledge in the early stages. Therefore, the knowledge of org. members gets reduced in the early stages of learning **[VII]**.

CHANGE IN PROPORTION OF INCORRECT KNOWLEDGE IN ORG. CODE, OVER TIME

Next, we look at the values of incorrect knowledge (wrong beliefs) percolating into the org. code for organizations with a moderate level of CHC (*Marchian* populations) in Figure A1C and for organizations with a low level of CHC (*sub-Marchian* populations) in Figure A1D. In Table T2, we present numerical values for the proportion of incorrect knowledge for the first six time-steps underlying these graphs. We note that the prediction made in [**IV**] and [**VI**] is borne out: the org. code contains a higher level of *incorrect* knowledge (upwards of 65%) for a *sub-Marchian* population (Figure A1C), compared to the level of *incorrect* knowledge in the org. code for a *Marchian* population (~40% and lower), as seen in Figure A1D.

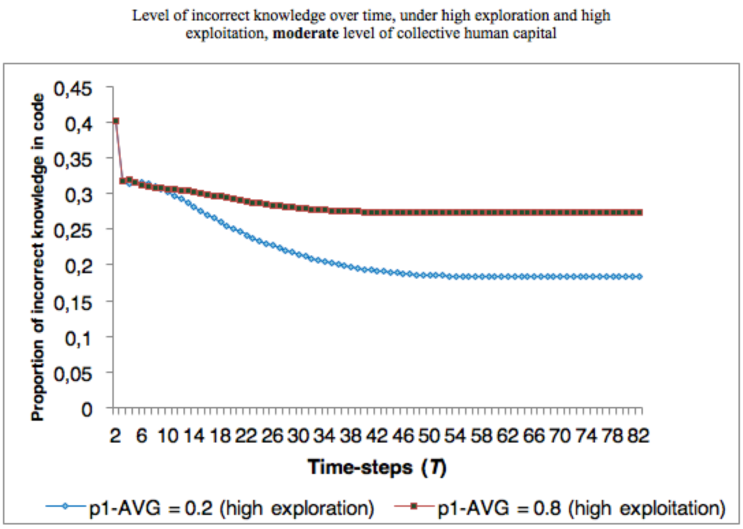

Figure A1C. Effect of heterogeneous learning rates ($p_{1\text{-AVG}}$) on the proportion of incorrect knowledge in the organizational code over time, *Marchian* populations

Notes. Rhombus-shaped markers represent members learning at an average rate of 20%, rectangular markers represent members learning at an average rate of 80%. In both cases, some members learn at a low rate (p_1 = 0.1), others learn at a high rate (p_1 = 0.9). Other parameters are M =30, N = 50, p_2 = 0.5, Iterations = 10,000, *Marchian* population.

24. We note that prediction [**III**] is borne out: the org. code contains a lower level of knowledge (~7%) for a sub-Marchian population (Lines L2 and L4, for t = 2, in Table T1), compared to org. code knowledge (~40%) for a Marchian population (Lines L1 and L3, for t = 2, in Table T1)

25. Though this analytical prediction reads similar to [**IV**], a fine point of difference is that the [**VI**] talks about inflow of incorrect knowledge into the org. code *after period two,* because the bar to becoming an elite is low.

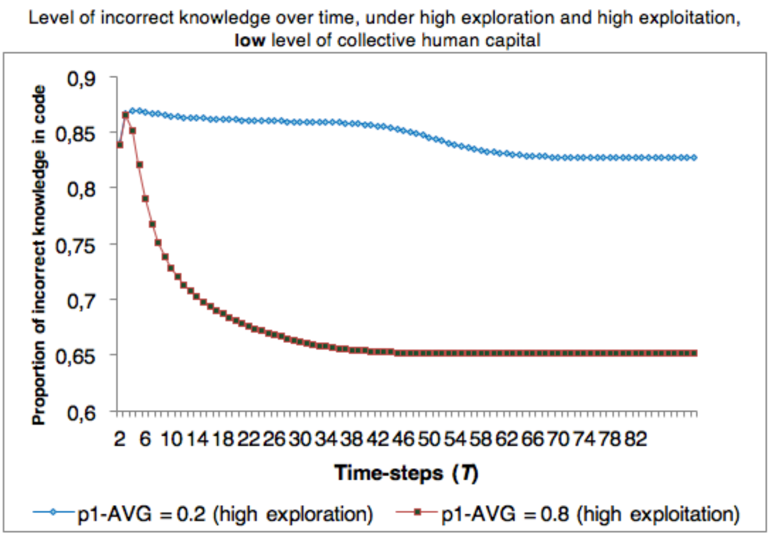

Notes. Rhombus-shaped markers represent members learning at an average rate of 20%, rectangular markers represent members learning at an average rate of 80%. In both cases, some members learn at a low rate (p_1 = 0.1), others learn at a high rate (p_1 = 0.9). Other parameters are M =30, N = 50, p_2 = 0.5, Iterations = 10,000, *sub-Marchian* population with deficiency 15% from *Marchian* population.

Figure A1D. Effect of heterogeneous learning rates ($p_{1\text{-AVG}}$) on the proportion of incorrect knowledge in the organizational code over time, *sub-Marchian* populations

Collective human capital		**Proportion of incorrect knowledge in org. code** High Exploration ($p_{1\text{-AVG}}$ = 0.2)						
Time-Step→		***t* = 1**	***t* = 2**	***t* = 3**		***t* = 4**	***t* = 5**	***t* = 6**
Moderate (*Marchian population*)	L1	0	0.402	0.317		0.314	0.316	0.315
Low (*sub-Marchian population*)	L2	0	0.839	0.867		0.869	0.869	0.868
		Proportion of incorrect knowledge in org. code High Exploitation ($p_{1\text{-AVG}}$ = 0.8)						
Time-Step→		***t* = 1**	***t* = 2**	***t* = 3**		***t* = 4**	***t* = 5**	***t* = 6**
Moderate (*Marchian population*)	L3	0	0.402	0.317		0.319	0.315	0.312
Low *sub-Marchian population*)	L4	0	0.838	0.866		0.851	0.821	0.791

Table T2: Proportion of **incorrect** knowledge in the org. code in the first six time-steps, for high exploration and high exploitation, for moderate collective human capital and low collective human capital

From lines L1 and L3 of Table T2 we learn that as in time-step *two* (t = 2)—when no member learning has yet taken place—the org. code houses about 40% incorrect knowledge, if CHC is moderate (*Marchian* populations). We further note that the percentage of incorrect knowledge comes down to about 31% rather quickly, i.e. in time-step t = 3. Thus, in this case, the probability of org. members obtaining false beliefs from the org. code is pretty low. We expect org. member knowledge to increase over time **[VIII]**. [26]

For low CHC (*sub-Marchian* populations), the org. code gets a high 84% wrong beliefs quickly by time-step *two* (t = 2) before any member learning has taken place, as shown by the values in lines L2 and L4. In this situation, we expect that the *Fast Learners* will imbibe false beliefs from the code, causing their correct knowledge to be reduced.[27] **[IX]**. The *Slow Learners* will also imbibe false beliefs from the org. code, but at a lower rate; their correct knowledge is also reduced, but at a lesser rate **[X]**.

SUMMARY: STATE OF ORG. CODE PRIOR TO EFFECTS OF MEMBER LEARNING

In Table T3 we present a summary of interpretations we make from the study of the state of the org. code at the end of time-step *two* (t = 2). We note that, before any member learning has taken place—i.e. before any exploration/exploitation activity has kicked in—the org. code attains very different character between *Marchian* and *sub-Marchian* populations (moderate and low levels of CHC, respectively). The low bar (~7%) for getting into the group of elites the org. code learns from, coupled with high incidence of incorrect knowledge in the org. code (84%) in *sub-Marchian* populations, presages the conclusion in proposition **2** that the level of org. knowledge developed will be much lower than that attained in the case of *Marchian* populations.

26. Though [**VIII**] reads similar to [**V**], a fine distinction is that [**VIII**] is premised on effective barring of incorrect beliefs lowering member knowledge onwards from step three.

27. We recall that March's model constructed a group of heterogeneous learners by having some members learn at a slow rate (p_1 = 0.1 for them), while others learn at a fast rate (p_1 = 0.9 for them). We designate the former group as *Slow Learners* and we designate the latter group as *Fast Learners*.

Initial Conditions	Time-step	Low CHC (*sub-Marchian* population)	Moderate CHC (*Marchian* population)
Org. code 0% correct knowledge, 0% incorrect knowledge	t = 0	**Org. members** 28% correct beliefs 43% incorrect beliefs	**Org. members** 33% correct beliefs 33% incorrect beliefs
LEADS TO: **Conditions just before start of member learning**			
	Time-step	*Low CHC*	*Moderate CHC*
State of org. code at end of time-step *two, prior to any member learning*	t = 2	**Org. code** 7% correct knowledge 84% incorrect knowledge	**Org. code** 40% correct knowledge 40% incorrect knowledge
INFERENCE	CREATING A SITUATION WHERE:		
At the time member learning kicks in, the character of the org. code is quite different for *Marchian* and *sub-Marchian* populations		The bar for getting into the elite group the org. code learns from is very low. Thus, more members with high number of incorrect beliefs advise the org. code	There is a high bar for getting into the elite group that the org. code learns from. The org. code is less likely to be advised by members with a high number of incorrect beliefs

Table T3: Summary interpretations from the state of the org. code in time-step *two* (t = 2)

EFFECT OF MEMBER LEARNING (EXPLORATION/ EXPLOITATION): MODERATE LEVEL OF CHC

In Figures A2A and A2B we present the average knowledge held by *Slow Learner* and *Fast Learner* organizational members (average member knowledge on the vertical axis) over time (time-steps on the horizontal axis), for high exploration ($\boldsymbol{p}_{1-AVG}$ = 0.2) and high exploitation ($\boldsymbol{p}_{1-AVG}$ = 0.8) respectively for an organization with a moderate level of CHC (a *Marchian* population).

Figures A2A and A2B inform us that for both high exploration ($\boldsymbol{p}_{1-AVG}$ = 0.2, Figure A2A) as well as high exploitation ($\boldsymbol{p}_{1-AVG}$ = 0.8, Figure A2B) the *Fast Learners* initially outpace the *Slow Learners* in acquiring new knowledge, as is to be expected.[28] The knowledge of *Fast Learners* increases faster because they are learning at a rate higher than the rate at which *Slow Learners* learn. Equilibrium is reached a few time-steps after the latter catch up with the former, as most of the organizational members increasingly carry correct knowledge (or properly aligned beliefs) in the same dimensions.[29] However, the time taken to catch up is higher in Figure A2A compared to that in Figure A2B. For example, in the case of high exploration ($\boldsymbol{p}_{1-AVG}$ = 0.2, Figure A2A), the difference in knowledge between the two groups of organizational members becomes less than 1% after time-step *fifty-five* (***t*** = 55). For high exploitation ($\boldsymbol{p}_{1-AVG}$ = 0.8, Figure A2B) this event occurs after time-step *forty-three* (***t*** = 43).

28. Also, note that the figures show that both classes of learners initiate with 33% correct knowledge, as we derived analytically in **[I]**, earlier. Further as predicted in **[V]** and **[VIII]**, knowledge of all org. members only increase over time.

29. Recall that the state of equilibrium is obtained when all the members attain an identical distribution of belief values, and this is reflected in the organizational code as well. Code learning stops because, at this point, the org. code fails to construct a group of elites it can learn from, given that no member carries more knowledge than the organizational code.

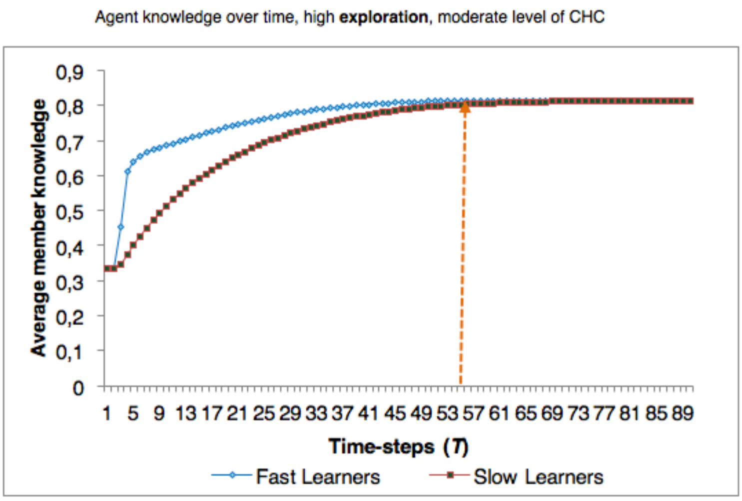

Notes. Rhombus-shaped markers represent *Fast Learners* who learn at a high rate (p_1 = 0.9). Rectangular markers represent *Slow Learners* who learn at a low rate (p_1 = 0.1). Other parameters are M =30, N = 50, p_{1-AVG} = 0.20, p_2 = 0.5, iterations = 10,000, *Marchian* population. The dotted arrow shows the point at which the two curves differ by less than 1% in value.

Figure A2A. Change in member knowledge over time, under high exploration, *Marchian* populations

Agent knowledge over time, high **exploitation**, moderate level of CHC

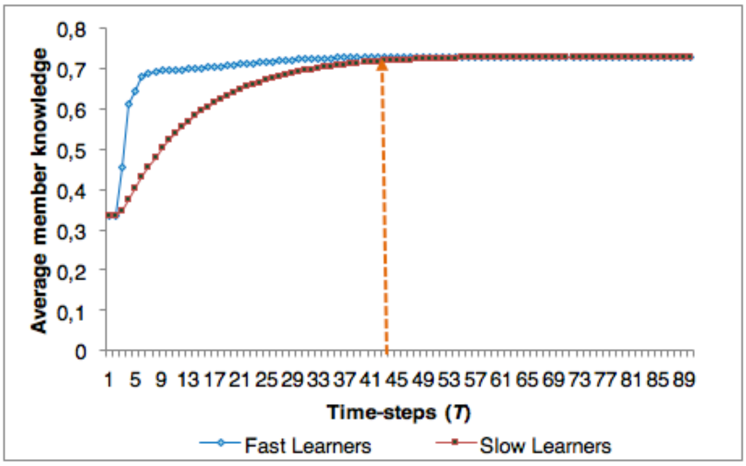

Notes. Rhombus-shaped markers represent *Fast Learners* who learn at a high rate (p_1 = 0.9). Rectangular markers represent *Slow Learners* who learn at a low rate (p_1 = 0.1). Other parameters are M =30, N = 50, p_{1-AVG} = 0.80, p_2 = 0.5, iterations = 10,000, *Marchian* population. The dotted arrow shows the point at which the two curves differ by less than 1% in value.

Figure A2B. Change in member knowledge over time, under high exploitation, *Marchian* populations

For the case of high exploration (p_{1-AVG} = 0.2) the catch-up (by *Slow Learners*) is delayed—compared to the case of high exploitation, (p_{1-AVG} = 0.8)—because a large number of organizational members are learning slowly. There is an important upside to a large number of organizational members learning slowly: a higher number of incorrect beliefs are flushed out of the org. code. This can be seen in Figure A1C. By around time-step *eighty*, the org. code "stabilizes" to having about 27% wrong beliefs for high exploitation (p_{1-AVG} = 0.8); the figure is around 18% for high exploration (p_{1-AVG} = 0.2). Why does the org. code stabilize with a higher number of incorrect beliefs when high exploitation is carried out by a *Marchian* population?

Recall (from lines L2 and L4 of Table T2, under t = 2) that the org. code (starting with zero incorrect beliefs and zero correct beliefs) imbibes 40% incorrect beliefs right after the first period of learning, before any org. member learning has taken place. When exploitation is high, a large number of *Fast Learners* will get some of these wrong beliefs from the org. code, including those who originally did not have a wrong belief in a particular knowledge dimension. When these members advise the org. code as elites in subsequent periods, the wrong beliefs in the org. code will persist. When exploration is high, a large number of members learn slowly. Fewer members will get wrong beliefs from the code. Correct beliefs will persist in other members. Later, when these members advise the org. code as members of the elite, some false knowledge gets flushed out from the org. code.

The persistence of a higher proportion of incorrect beliefs for high exploitation turns out to be decisive: higher exploration yields higher organizational knowledge compared to higher exploitation. Thus, when the CHC in an organization is moderate (as embodied in a *Marchian* population) or higher, high exploration is preferable over high exploitation. Proposition **1** is substantiated.

Additional analysis, based on learning by the org. code. Figure A2C (high exploration, p_{1-AVG} = 0.2) and Figure A2D (high exploitation, p_{1-AVG} = 0.8) reveal further details as to why, for *Marchian* populations, higher exploration leads to higher overall organizational knowledge. In both these figures, time-steps of simulation are shown on the horizontal axis. On the vertical axis, we show the net number of knowledge dimensions for which organizational elites provide *good recommendations* (i.e. recommendations that are in line with the requirements of the reality **R**) and the number of knowledge dimensions for which the organizational elites provide *bad recommendations* (i.e. the elite recommendations are wrong). To compute the values of the parameters shown on the vertical axis, we ignore the cases where the elites provide a correct recommendation for a knowledge dimension that the organization code already knows correctly. We also ignore the cases where the elites provide a wrong recommendation for a knowledge dimension for which the org. code already has incorrect knowledge.

We observe that *net good recommendations* are available for longer and in greater bulk in Figure A2C (high exploration, moderate CHC) compared to in Figure A2D (high exploitation, moderate CHC). This is the implication of March's statement that the presence of a higher proportion of *Fast Learners* leads to quicker degeneration of knowledge heterogeneity in the organization. Thus exploration (in companies with moderate CHC) attains outcomes superior to exploitation because good quality recommendations from the elite are available for a longer time

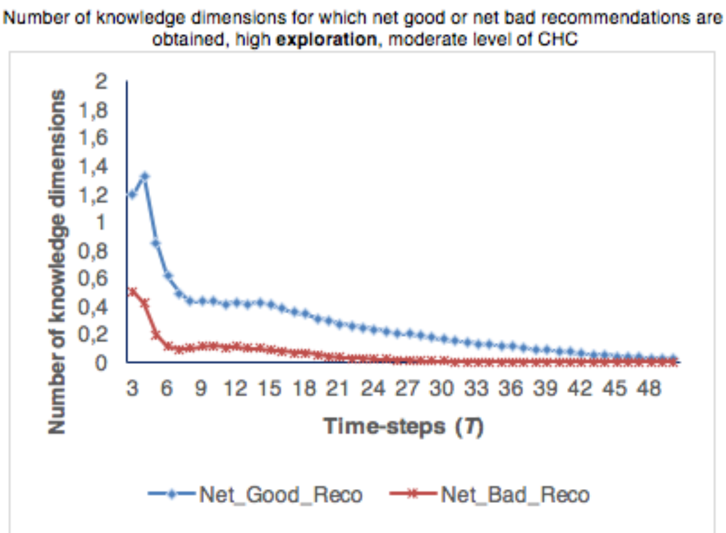

Notes. Rhombus-shaped markers represent net number of knowledge dimensions for which elites recommend correctly. Rectangular markers represent net number of knowledge dimensions for which elites recommend wrongly. Other parameters are ***M*** =30, ***N*** = 50, $\boldsymbol{p}_{1\text{-AVG}}$ = 0.20, $\boldsymbol{p}_2$ = 0.5, Iterations = 10,000, *Marchian* population.

Figure A2C. Net good and bad recommendations from the elite over time, under high exploration, *Marchian p*opulations

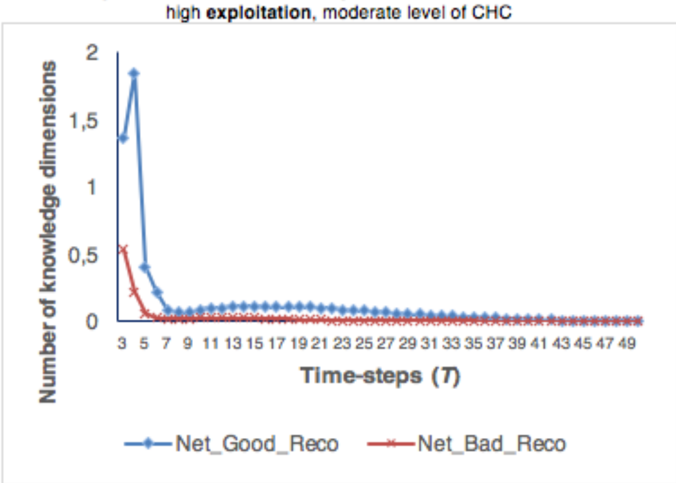

Notes. Rhombus-shaped markers represent net number of knowledge dimensions for which elites recommend correctly. Rectangular markers represent net number of knowledge dimensions for which elites recommend wrongly. Other parameters are M =30, N = 50, p_{1-AVG} = 0.80, p_2 = 0.5, Iterations = 10,000, *Marchian* population.

Figure A2D. Net good and bad recommendations from the elite over time, under high exploitation, *Marchian* populations

EFFECT OF MEMBER LEARNING—LOW LEVEL OF CHC

In Figures A3A and A3B we present average knowledge held by *Slow Learner* and *Fast Learner* members of the organization (average member knowledge on the vertical axis) over time (time-steps on the horizontal axis), for high exploration (p_{1-AVG} = 0.2) and high exploitation (p_{1-AVG} = 0.8) respectively for *sub-Marchian* populations. As predicted in [**II**] all org. members start with about 28% knowledge. Moreover, as predicted in [**IX**], the knowledge of the *Fast Learners* reduces in the initial stages, since they learn from a highly faulty org. code. Further, as predicted in [**X**], the knowledge of the *Slow Learners* also goes down in the initial stages for the same reason; for them the rate of fall is less steep, since they learn faulty information at a slower rate.

In Table T4 we show knowledge in org. code, and that in *Slow* and *Fast Learners*, and proportion of incorrect knowledge in the org. code, for high exploration and high exploitation, for a *sub-Marchian* population (i.e. where the organization possesses low CHC). The numbers presented in lines L2 and L3 comprise the values underlying the graphs shown in Figure A3A. The numbers presented in lines L6 and L7 comprise the values underlying the graphs shown in Figure A3B. Lines L1 and L5 are imported from Table T1. Lines L4 and L8 are imported from Table T2.

Low CHC		**Knowledge level for** High **Exploration** (p_{1-AVG} = 0.2)					
sub-Marchian population		***t*** **= 1**	***t*** **= 2**	***t*** **= 3**	***t*** **= 4**	***t*** **= 5**	***t*** **= 6**
Org. code knowledge (*T1-L2*)	L1	0	0.069	0.086	0.104	0.117	0.125
Knowledge of *Fast Learners*	L2	0.283	0.283	0.118	0.103	0.112	0.121
Knowledge of *Slow Learners*	L3	0.283	0.283	0.265	0.249	0.235	0.223
Proportion of incorrect knowledge in org. code (*T2-L2*)	L4	0	0.839	0.867	0.869	0.869	0.868
		Knowledge level for High **Exploitation** (p_{1-AVG} = 0.8)					
		t **= 1**	***t*** **= 2**	***t*** **= 3**	***t*** **= 4**	***t*** **= 5**	***t*** **= 6**
Org. code knowledge (*T1-L4*)	L5	0	0.069	0.086	0.139	0.176	0.209
Knowledge of *Fast Learners*	L6	0.283	0.283	0.120	0.105	0.138	0.172
Knowledge of *Slow Learners*	L7	0.283	0.283	0.265	0.249	0.238	0.232
Proportion of Incorrect knowledge in org. code (*T2-L4*)	L8	0	0.838	0.866	0.851	0.821	0.791

Please note: A label such as T*x*-L*y* designates the table number and line number of the source table. For example, the last row displays a label T2-L4. This means that the information is sourced from line four of table T2.

Table T4: Proportion of correct knowledge in the organizational code, in *Fast Learners* and *Slow Learners*, and proportion of incorrect knowledge in the org. code, in the first six time-steps, for high exploration and high exploitation, for low CHC (*sub-Marchian* population)

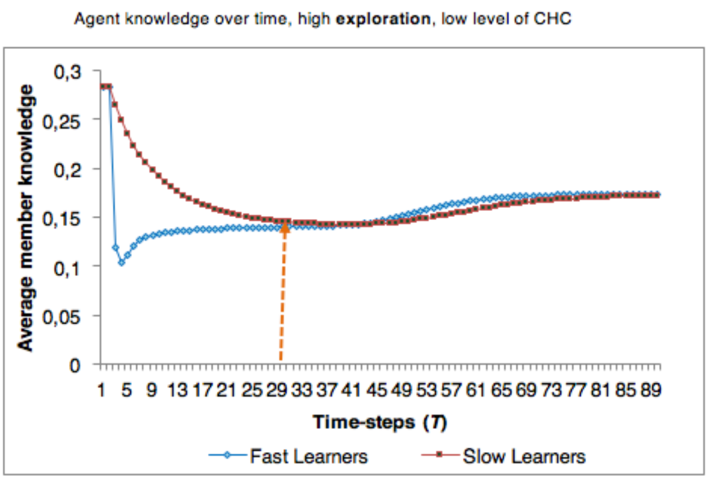

Notes. Rhombus-shaped markers represent *Fast Learners* who learn at a high rate (p_1 = 0.9). Rectangular markers represent *Slow Learners* who learn at a low rate (p_1 = 0.1). Other parameters are M =30, N = 50, p_{1-AVG} = 0.20, p_2 = 0.5, iterations = 10,000, *sub-Marchian* population with deficiency 15% from *Marchian* population. The dotted arrow shows the point at which the two curves differ by less than 1% in value at the time the *Fast Learners* catch up with *Slow Learners*.

Figure A3A. Effect of high exploration on member knowledge over time, *sub-Marchian* populations

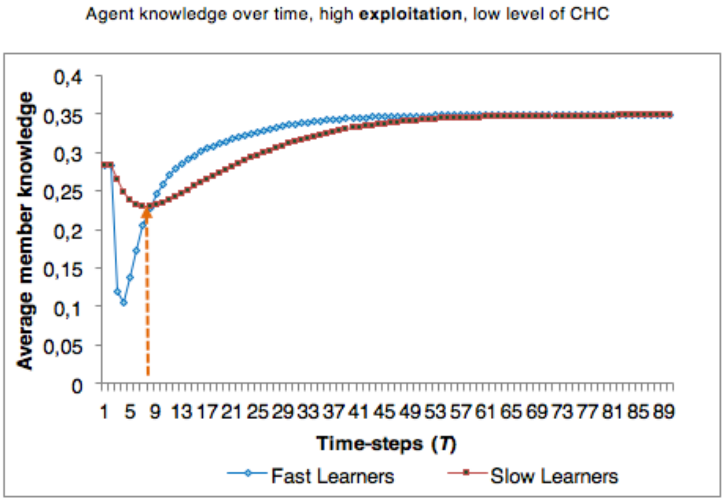

Notes. Rhombus-shaped markers represent *Fast Learners* who learn at a high rate (p_1 = 0.9). Rectangular markers represent *Slow Learners* who learn at a low rate ($\boldsymbol{p_1}$ = 0.1). Other parameters are $\boldsymbol{M}$ =30, $\boldsymbol{N}$ = 50, $\boldsymbol{p_{1-AVG}}$ = 0.80, $\boldsymbol{p_2}$ = 0.5, iterations = 10,000, *sub-Marchian* population with deficiency 15% from *Marchian* population. The dotted arrow shows the point at which the two curves differ by less than 1% in value at the time the *Fast Learners* catch up with *Slow Learners.*

Figure A3B. Effect of high exploitation on member knowledge over time, *sub-Marchian* populations

The phenomenon of org. member knowledge in *sub-Marchian* populations decreasing in the first few time-steps (as predicted in [**VII**]) is the opposite of what happens in *Marchian* populations (where, as we have previously analytically argued and graphically demonstrated, member knowledge goes up till equilibrium is reached). Let us see why this happens.

In *sub-Marchian* populations, an early increase in misaligned beliefs in *Fast Learners* occurs owing to contamination from vast numbers of incorrect knowledge bits in the org. code. We recall that early on (time-step *two*, $\boldsymbol{t}$ = 2) the org. code gets populated with a very high proportion of incorrect knowledge when the organizational population is *sub-Marchian* (84% as seen in Figure A1D). The knowledge level of *Fast Learners* in *sub-Marchian* populations goes down from the 28% initial level ($\boldsymbol{t}$ = 0) to about a 10% level at the lowest point ($\boldsymbol{t}$ = 4) because *Fast Learners* pick incorrect knowledge from the org. code. Further, as analytically reasoned earlier, the *Slow Learners* suffer less dilution of knowledge on account of learning slowly.

In *sub-Marchian* populations, for both exploration (Figure A3A) and exploitation (Figure A3B), the org. code knowledge keeps increasing (from the initial level of 7% at the end of time-step *two*, $\boldsymbol{t}$ = 2) after time-step *two* ($\boldsymbol{t}$ = 2) because org. members—particularly the *Slow Learners*—have more knowledge than the org. code, enabling formation of the elite group that the org. code learns from. Further, the accumulation of incorrect knowledge by *Fast Learners* eventually stops at about 10% value, and thereafter the knowledge of *Fast Learner* members starts increasing. The knowledge of *Fast Learners* does not go down below 10% because, by around time-step *four* ($\boldsymbol{t}$ = 4), the org. code knowledge (that was about 7% in time-step *two,* $\boldsymbol{t}$ = 2) increases to a value higher than 10% by extracting some correct knowledge from the diverse knowledge of *Slow Learners* in the interim period.[30]

We notice that, both in Figure A3A and Figure A3B, till the knowledge of the Fast Learners catches up with the knowledge of Slow Learners, the knowledge of the Slow Learners continues to be eroded on account of socialization into wrong beliefs of the (highly faulty) org. code.

30. This can be seen in Table T1, column $\boldsymbol{t}$ = 4, from line L2 for the high exploration case (org. code knowledge = 10.2%), and line L4 for the high exploitation case (org. code knowledge = 13.9%).

Interestingly, the catch-up to Slow Learner knowledge level by Fast Learners happens at the 15% mark at around time-step thirty for high exploration (Figure A3A), and at above 25% for high exploitation (Figure A3B) by time-step eight. As expected, when the organization is composed of a large proportion of Fast Learners, catch-up happens earlier and when the organization is composed of a large proportion of Slow Learners, catch-up happens late.

The catch-up event is important. Before catch-up, a *Slow Learner's* unique correct knowledge is under a strong threat of being updated by faulty knowledge from the org. code. At the time of catch-up, *Fast Learners* carry correct knowledge in largely the same dimensions for which the org. code carries correct knowledge. Therefore, *Fast Learners* carrying lots of faulty knowledge will not make it to the team of elites that advise the org. code after catch-up. When the org. code receives less bad advice, there is a lower probability that a *Slow Learner* will be injected with wrong knowledge from the org. code. Post catch-up, when a *Slow Learner* carrying unique correct knowledge makes it to the team of elites advising the org. code, the advice to the org. code is increasingly regarding the unique correct knowledge dimensions about which the org. code is unaware (i.e. the org. code carries faulty knowledge or neutral beliefs).

Thus, under high exploitation, because the knowledge of the *Slow Learners* is eroded for a shorter time, the catch-up happens when *Slow Learners'* knowledge still has a relatively high value, 25%, compared to the case of high exploration, where the catch-up happens much later, by which time the *Slow Learners'* knowledge has already fallen to 15%. In either case, thereafter, the org. code knowledge rises somewhat, as the final bits of heterogeneous knowledge from the org. members are rounded up as org. code knowledge.

Additional analysis, based on learning by the organizational code. As in the previous case, we may also take a look under the hood by considering the quality of recommendations elites make, when exploration or exploitation is carried out by *sub-Marchian* populations. In Figure A3C, we show the numbers of (net) good and bad recommendations made by the elites, for high exploration. In Figure A3D we show the corresponding picture for high exploitation. On the vertical axis, we show the number of knowledge dimensions for which the elites have a good recommendation (i.e. the recommendation from the elite matches the value in external reality **R**) and the number of knowledge dimensions for which the elite have a bad recommendation (i.e. elite recommendation does not match the value in external reality **R**). As before, we ignore (a) cases where elites have made a good recommendation but the org. code already has correct knowledge, and (b) cases where elites make an incorrect recommendation for a dimension on which the org. code knowledge is already wrong. We observe that good recommendations from elites are less voluminous and degenerate rapidly for high exploration (Figure A3C) compared to high exploitation (Figure A3D). *The copious good recommendations early on—as observed in Figure A3D—help organizational knowledge rise faster, enabling Fast Learners to catch up Slow Learners before the latter degenerate significantly.*[31] A comparison of lines L1 (high exploration) and L5 (high exploitation) in Table T4 shows this. In both, we find that org. code knowledge is about 8.6% at t =3. For L1 (high exploration) org code knowledge rises slowly recording 10.4%, 11.7% and 12.5% in the next three time-steps. The corresponding figures are 13.9%, 17.6% and 20.9% for L5 (high exploitation).

31. *Fast Learners* quickly learn from the fast-rising org. code knowledge.

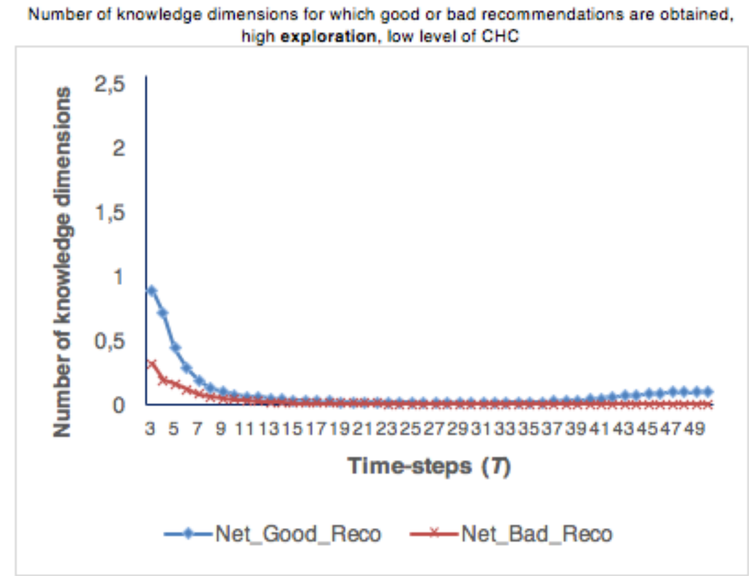

Notes. Rhombus-shaped markers represent number of knowledge dimensions for which elites recommend correctly. Rectangular markers represent number of knowledge dimensions for which elites recommend wrongly. Other parameters are ***M*** =30, ***N*** = 50, $\boldsymbol{p}_{1\text{-AVG}}$ = 0.20, $\boldsymbol{p}_2$ = 0.5, Iterations = 10,000, *Marchian* population.

Figure A3C. Net good and bad recommendations from the elite over time, under high exploration, *sub-Marchian* populations

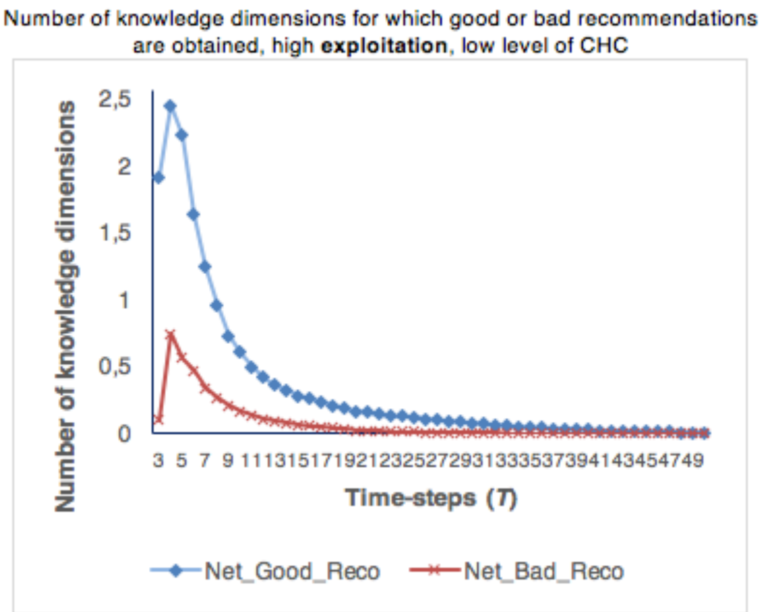

Notes. Rhombus-shaped markers represent number of knowledge dimensions for which elites recommend correctly. Rectangular markers represent number of knowledge dimensions for which elites recommend wrongly. Other parameters are ***M*** =30, ***N*** = 50, $\boldsymbol{p}_{1\text{-AVG}}$ = 0.80, $\boldsymbol{p}_2$ = 0.5, Iterations = 10,000, *Marchian* population.

Figure A3D. Net good and bad recommendations from the elite over time, under high exploitation, *sub-Marchian* populations

SUMMARY: WHY DOES MODEL BEHAVIOR CHANGE UPON CHANGE OF CHC?

In table **T5**, we provide a summary of the description and reasoning provided above.

Sl.	Moderate CHC (*Marchian* population)	Low CHC (*Sub-Marchian* population)
1.	Before any member learning occurs, the org. code attains ~40% correct knowledge and ~40% incorrect knowledge, by end of period 2.	Before any member learning occurs, the org. code attains ~7% correct knowledge and ~84% incorrect knowledge, by end of period 2.
2.	Because the bar for getting into elites is high (~40%), member knowledge can only increase (from the initial level of 33%) upon learning from the org. code. Knowledge of *Fast Learners* rises faster.	Because the bar for getting into elites is low (~7%), member knowledge decreases in the early time-steps, from the initial level of 28%. Knowledge of *Slow Learners* decreases slower.
3.	*Slow Learners* will have to catch up with the knowledge level of *Fast Learners* before the org. code reaches the equilibrium knowledge level.	*Fast Learners* will have to catch up with the knowledge level of *Slow Learners* before the Org Code reaches the equilibrium knowledge level.
4.	Catch-up happens earlier when the proportion of *Fast Learners* is high (i.e. under high exploitation).	Catch-up happens earlier when the proportion of *Fast Learners* is high (i.e. under high exploitation).
5.	Exploration benefits transpire through a higher proportion of *Slow Learners*, owing to knowledge heterogeneity lasting for longer (compared to the duration that knowledge heterogeneity lasts in the case of high exploitation).	Exploitation benefits transpire through having a higher proportion of *Fast Learners*, by halting the loss of knowledge heterogeneity of *Slow Learners* earlier (compared to the duration that knowledge heterogeneity is lost, in the case of high exploration).
6.	Much more heterogeneity is utilized, enriching the org. code, owing to org. members being exposed to low levels of faulty knowledge. Incorrect knowledge in the org. code falls from ~40% in time-step *two* to ~31% by time-step *three*. Final org. code knowledge is upwards of 70% (Figure A1A).	Much less member heterogeneity is utilized, leading to final org. code knowledge being 35% or lower (Figure A1B). This is because members are exposed to a high extent of faulty knowledge in the org. code. Org. code has upwards of 86% incorrect knowledge at time-step *three*.

Table T5: Summary of the micro-level dynamics under low and moderate CHC

APPENDIX B

Variation of critical level of collective human capital (CHC) with organizational size and breadth or context

In Figure B1 we show the variation of the critical value of CHC at which a desirable strategy of high exploitation gives way to a strategy favoring high exploration when the size of an organization (number of members, ***N***) is increased in a stable environment. On the vertical axis, we plot CHC deficiency at the critical point—where the curves $p_{1\text{-AVG}} = 0.2$ and $p_{1\text{-AVG}} = 0.8$ intersect (as shown in Figure 5)—as per cent offset from the CHC of a *Marchian* population. Figure B2 provides analogous information for variation in the number of dimensions of the environment or reality (breadth of context, ***M***). Figures B3 and B4 provide corresponding information for a dynamic environment having 2% turbulence. In all the figures, we observe that the critical value of CHC increases and approaches that of a *Marchian* population upon an increase in values of the parameter on the horizontal axis (size of organization/breadth of context).

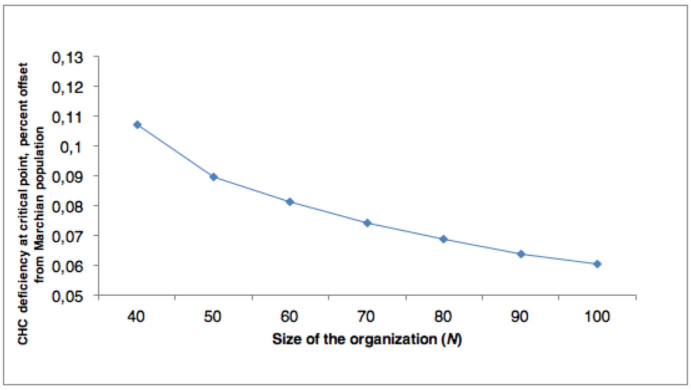

Parameters. M =30, p_2 = 0.5, p_4 = 0, Iterations = 10,000
Figure B1. Effect of size of organization on location of critical point, stable environment

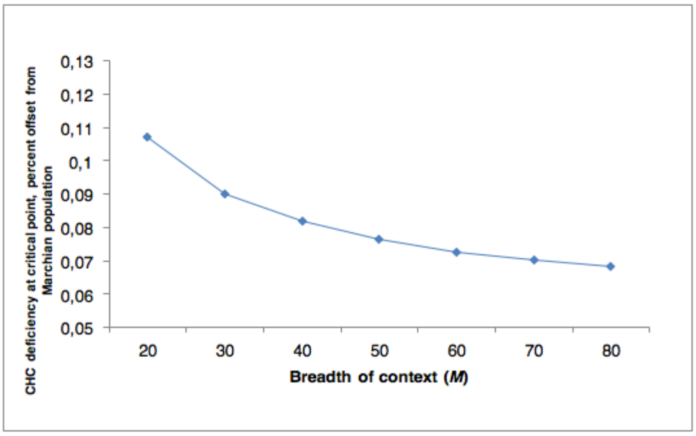

Parameters. N =50, p_2 = 0.5, p_4 = 0, Iterations = 10,000.
Figure B2. Effect of breadth of context on location of critical point, stable environment

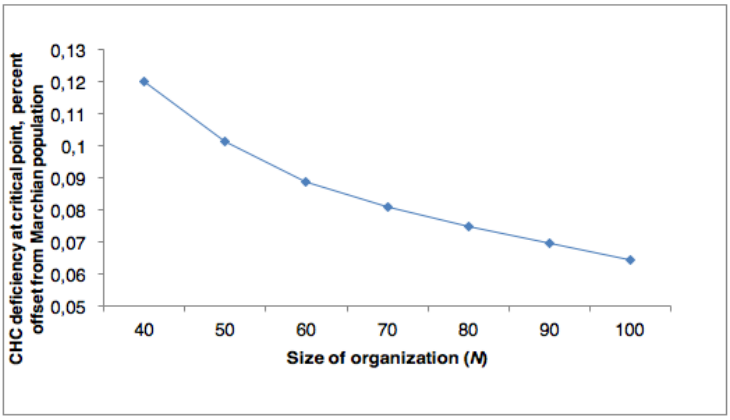

Parameters. M =30, p_2 = 0.5, p_4 = 0.02, Iterations = 10,000.
Figure B3. Effect of size of organization on location of critical point, dynamic environment

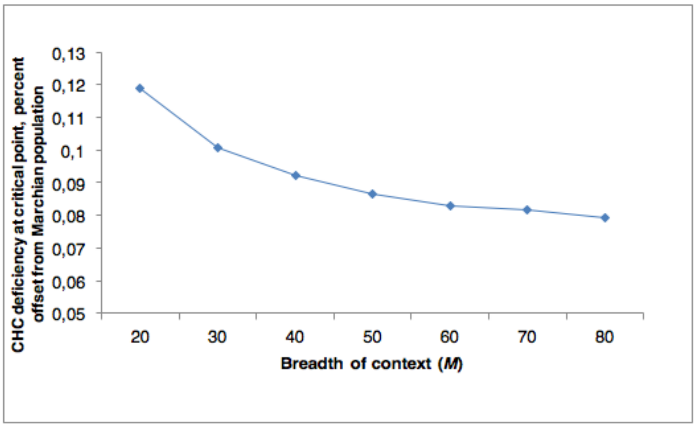

Parameters. N =50, p_2 = 0.5, p_4 = 0.02, Iterations = 10,000.
Figure B4. Effect of breadth of context on location of critical point, dynamic environment

REFERENCES

Adler, P.S., Goldoftas, B. & Levine, D.I. (1999). Flexibility versus Efficiency? A Case Study of Model Changeovers in the Toyota Production System. *Organization Science*, 10(1), 43-68.

Adner, R., Pólos, L., Ryall, M. & Sorenson, O. (2009). The Case for Formal Theory. *Academy of Management Review*, 34(2), 201-208.

Becker, G.S. (1964), *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education,* New York, NY: Columbia University Press.

Bendor, J., Moe, T.M. & Shotts, K.W. (2001). Recycling the Garbage Can: An Assessment of the Research Program. *American Political Science Review*, 95(1), 169-190.

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

Birkinshaw, J., Zimmerman, A. & Raisch, S. (2016). How Do Firms Adapt to Discontinuous Change? Bridging the Dynamic Capabilities and Ambidexterity Perspectives. *California Management Review*, 58(4): 36-58.

Blaschke, S. & Schoeneborn, D. (2006). The Forgotten Function of Forgetting: Revisiting Exploration and Exploitation in Organizational Learning. *Soziale Systeme*, 11(2), 99-119.

Brady, T. & Davies, A. (2004). Building Project Capabilities: From Exploratory to Exploitative Learning. *Organization Studies*, 25(9), 1601-1621.

Chanda, S.S. (2017). Inferring Final Organizational Outcomes from Intermediate Outcomes of Exploration and Exploitation: The Complexity Link. *Computational and Mathematical Organization Theory*, 23(1), 61-93.

Chanda, S.S. & Ray, S. (2015). Optimal Exploration and Exploitation: The Managerial Intentionality Perspective. *Computational and Mathematical Organization Theory*, 21(3), 247-273.

Chen, E.L. & Katila, R. (2008). Rival Interpretations of Balancing Exploration and Exploitation: Simultaneous or Sequential? In S. Shane, (Ed.), *Handbook of Technology and Innovation Management* (pp. 197-214). New York, NY: John Wiley and Sons Ltd.

Child, J. (1972). Organizational Structure, Environment and Performance: The Role of Strategic Choice. *Sociology*, 6(1), 1-22.

Dierickx, I. & Cool, K. (1989). Asset Accumulation and Sustainability of Competitive Advantage. *Management Science*, 35(12), 1504-1511.

Duncan, R.B. (1972). Characteristics of Organizational Environments and Perceived Environmental Uncertainty. *Administrative Science Quarterly*, 17(3), 313-327.

Eisenhardt, K.M. & Brown S.L. (1997). The Art of Continuous Change: Linking Complexity Theory and Time-paced Evolution in Relentlessly Shifting Organizations. *Administrative Science Quarterly*, 42(1), 1–34.

Fang, C., Lee, J. & Schilling, M.A. (2010). Balancing Exploration and Exploitation through Structural Design: The Isolation of Subgroups and Organizational Learning. *Organization Science*, 21(3), 625-642.

Farjoun, M. (2010). Beyond Dualism: Stability and Change as Duality. *Academy of Management Review*, 35(2), 202-225.

Fillenbaum, S. (1974). Information Amplified: Memory for Counterfactual Conditionals. *Journal of Experimental Psychychology*, 102(1), 44-49.

Garcias, F., Dalmasso, C. & Sardas, J.C. (2015). Paradoxical Tensions in Learning Processes: Exploration, Exploitation and Exploitative Learning. *M@n@gement*, 18(2), 156-178.

Gavetti, G, Greve, H.R., Levinthal, H.A. & Ocasio, W. (2012). The Behavioral Theory of the Firm: Assessment and Prospects. *Academy of Management Annals*, 6(1), 1-40.

Gibson, C.B. & Birkinshaw, J. (2004). The Antecedents, Consequences, and Mediating Role of Organizational Ambidexterity. *Academy of Management Journal*, 47(2), 209-226.

Gupta, A., Smith, K. & Shalley, C. (2006). The Interplay between Exploration and Exploitation. *Academy of Management Journal*, 49(4), 693-706.

Hamel, G. & Prahalad, C.K. (1993). Strategy as Stretch and Leverage. *Harvard Business Review*, March-April, 75-84.

Hargadon, A. & Fanelli, A. (2002). Action and Possibility: Reconciling Dual Perspectives of Knowledge in Organizations. *Organization Science*, 13(3), 290-302.

Holland, J. (1975), *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control & Artificial Intelligence,* Ann Arbor, MI: University of Michigan Press.

Huber, G.P. (1991). Organizational Learning: The Contributing Processes and the Literature. *Organization Science*, 2(1), 88-115.

Hutzschenreuter, T., Pedersen, T. & Volberda, H.W. (2007). Internationalization—the Role of Managerial Intentionality and Path Dependency: A Perspective on International Business Research. *Journal of International Business Studies*, 38(7), 1055-1068.

ICMR (2001). The Nirma Story. *IBS Center for Management Research*. Available at: http://www.icmrindia.org/casestudies/catalogue/Marketing/MKTG008.htm (accessed on March 10, 2016).

Jansen, J.J.P., Van den Bosch, F.A.J. & Volberda. H.W. (2006). Exploratory Innovation, Exploitative Innovation, and Performance: Effects of Organizational Antecedents and Environmental Moderators. *Management Science*, 52(11), 1661-1674.

Kane, G.C. & Alavi, M. (2007). Information Technology and Organizational Learning: An Investigation of Exploration and Exploitation Processes. *Organization Science*, 18(5), 796-812.

Kim, T. & Rhee, M. (2009). Exploration and Exploitation: Internal Variety and Environmental Dynamism. *Strategic Organization,* 7(1), 11-41.

Kogut, B. & Zander, U. (1992). Knowledge of the Firm, Combination Capabilities, and the Replication of Technology. *Organization Science*, 3(3), 383-397.

Kogut, B. & Zander, U. (1996). What Firms Do? Coordination, Identity and Learning. *Organization Science*, 7(5), 502-518.

Kuran, T. (1988). The Tenacious Past: Theories of Personal and Collective Conservatism. *Journal of Economic Behavior and Organization*, 10(2), 143-171.

Lavie, D., Kang, J. & Rosenkopf, L. (2011). Balance within and across Domains: The Performance Implications of Exploration and Exploitation in Alliances. *Organization Science*, 22(6), 1517-1538.

Lazer, D. & Friedman, A. (2007). The Network Structure of Exploration and Exploitation. *Administrative Science Quarterly*, 52(4), 667-694.

Lee, J.Y., Bachrach, D.G. & Rousseau, D.M. (2015). Internal Labor Markets, Firm-specific Human Capital, and Heterogeneity Antecedents of Employee Idiosyncratic Deal Requests. *Organization Science*, 26(3), 794-810.

Levinthal, D.A. & March, J.G. (1993). The Myopia of Learning. *Strategic Management Journal*, Winter Special Issue 14(S2), 95-112.

Levitt, B. & March J.G. (1988). Organizational Learning. *Annual Review of Sociology*, 14(1), 319-338.

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

March, J.G. (2013). In Praise of Beauty. *M@n@gement,* 16(5), 732-738.

Miller, K.D. & Martignoni, D. (2016). Organizational Learning with Forgetting: Reconsidering the Exploration–Exploitation Tradeoff. *Strategic Organization*, 14(1), 53-72.

Miller, K.D., Zhao, M. & Calantone, R.J. (2006). Adding Interpersonal Learning and Tacit Knowledge to March's Exploration-Exploitation Model. *Academy of Management Journal*, 49(4), 709-722.

Nelson, R.R. & Winter, S.G. (1982), *An Evolutionary Theory of Economic Change,* Cambridge, MA: The Belknap Press of Harvard University Press.

Nonaka, I. & Takeuchi, H. (1995), *The Knowledge Creating Company,* New York, NY: Oxford University Press.

Nonaka, I. & von Krogh, G. (2009). Tacit Knowledge and Knowledge Conversion: Controversy and Advancement in Organizational Knowledge Creation Theory. *Organization Science*, 20(3), 635–652.

Pearl, J. (2000), *Causality: Models, Reasoning, and Inference,* New York, NY: Cambridge University Press.

Ployhart, R.E. & Moliterno, T.P. (2011). Emergence of the Human Capital Resource: A Multilevel Model. *Academy of Management Review*, 36(1), 127-150.

Posen, H.E. & Levinthal, D.A. (2012). Chasing a Moving Target: Exploitation and Exploration in Dynamic Environments. *Management Science*, 58(3), 587-601.

Prahalad, C.K. (2004), *The Fortune at the Bottom of the Pyramid: Eradicating Poverty through Profits*, Philadelphia, PA: Wharton School Publishing.

Prahalad, C.K. & Hart, S.L. (2002). The Fortune at the Bottom of the Pyramid. *strategy+business*, 26(1), 1-14.

Rodan, S. (2005). Exploration and Exploitation Revisited: Extending March's Model of Mutual Learning. *Scandinavian Journal of Management*, 21(4), 407-428.

Schilling, M.A. & Fang, C. (2014). When Hubs Forget, Lie, And Play Favorites: Interpersonal Network Structure, Information Distortion, and Organizational Learning. *Strategic Management Journal*, 35(7), 974-994.

Schmitt, A., Probst, G. & Tushman, M.L. (2010). M@n@gement in Times of Economic Crisis: Insights into Organizational Ambidexterity. *M@n@gement*, 13(3), 128-150.

Schumpeter, J.A. (1934), *The Theory of Economic Development*. Cambridge, MA: Harvard University Press.

Siggelkow, N. & Rivkin, J.W. (2006). When Exploration Backfires: Unintended Consequences of Multilevel Organizational Search. *Academy of Management Journal*, 49(4), 779-795.

Simon, H.A. (1991). Bounded Rationality and Organizational Learning. *Organization Science*, 2(1), 125-134.

Smith, K.G., Collins, C.J. & Clark, K.D. (2005). Existing Knowledge, Knowledge Creation Capability, and the Rate of New Product Introduction in High-technology Firms. *Academy of Management Journal*, 48(2), 346-357.

Smith, W.K., Binns, A. & Tushman, M.L. (2010). Complex Business Models: Managing Strategic Paradoxes Simultaneously. *Long Range Planning*, 43(2), 448-461.

Spender, J.C. (2008). Organizational Learning and Knowledge Management: Whence and Whither? *Management Learning*, 39(2), 159-176.

Stettner, U. & Lavie, D. (2014). Ambidexterity under Scrutiny: Exploration and Exploitation via Internal Organization, Alliances, and Acquisitions. *Strategic Management Journal*, 35(13), 1903-1929.

Tushman, M.L. & O'Reilly, C.A. (1996). Ambidextrous Organizations: Managing Evolutionary and Revolutionary Change. *California Management Review*, 38(4), 8-30.

von Nordenflycht, A. (2011). Firm Size and Industry Structure under Human Capital Intensity: Insights from the Evolution of the Global Advertising Industry. *Organization Science*, 22(1), 141-157.

Zhang, H. & Xi, Y. (2010). Exploration and Exploitation in Parallel Problem Solving: Effect of Imitation Strategy and Network Structure. *International Journal of Knowledge and Systems Science*, 1(3), 55-67.

Sasanka Sekhar Chanda is Associate Professor in Strategic Management in the Indian Institute of Management Indore. His research interests are in strategic decision-making, managerial intentionality, and complexity theory. Previously, he worked for thirteen years in industry in a range of roles spanning engineering, consulting and management.

Sougata Ray is Professor in Strategic Management in the Indian Institute of Management Calcutta. He has been researching for nearly two decades on how policy-induced changes in institutional environment influence the strategic behavior of emerging economy firms and how differences in quasi institutional environment, heterogeneity of firm characteristics and strategic intent of the top management lead to variations in strategy for growth, internationalization, innovations, entrepreneurship and sustainability of firms, business groups and state-owned enterprises.

Bill Mckelvey is Professor Emeritus of Strategic Organizing, Complexity Science and Econophysics at UCLA Anderson. Early articles focus on organization and socio-technical systems design. His book, *Organizational Systematics* (1982), remains the definitive treatment of organizational taxonomy and evolutionary theory. As Director of the Center for Rescuing Strategy & Organization Science (SOS), he initiated activities leading to the founding of UCLA's Center for Human Complex Systems and Computational Social Science.

Acknowledgments : We express our deep gratitude to Dr Thomas Roulet and three anonymous reviewers for significant help and guidance in improving the manuscript. All errors remain the authors' sole responsibility.