

STRATEGIC EXPLOITATION CHALLENGES AND ORGANIZATIONAL LEARNING: INTEGRATING COMPUTATIONAL AND FIELD RESEARCH FINDINGS

Robert A Burgelman
Edmund W. Littlefield Professor of Management
Graduate School of Business
Stanford University
655 Knight Way, Stanford, CA 94305, U.S.A.
Phone Number: 1 (650) 723-4488
profrab@stanford.edu

Sasanka Sekhar Chanda*
Professor, Strategic Management, IIM Indore
C-101 Academic Block,
Indian Institute of Management Indore
Prabandh Shikhar, Rau-Pithampur Road,
Indore, Madhya Pradesh 453556, India
Phone number: +917312439591
sschanda@iimidr.ac.in

April 8, 2026

[* Corresponding Author]

Acknowledgement

We dedicate this paper to the living memory of Jim March's inspiring scholarship.

STRATEGIC EXPLOITATION CHALLENGES AND ORGANIZATIONAL LEARNING: INTEGRATING COMPUTATIONAL AND FIELD RESEARCH FINDINGS

Abstract

We report computational analysis with extensions of March's (1991) genetic model of exploration and exploitation to further illuminate strategic exploitation challenges highlighted in field research about the transition from exploration to exploitation in strategic innovation projects. We disentangle the effects, under different conditions of *exploitation urgency*, of different conjunctions of *exploitation drive* and *exploitation capability* on organizational learning outcomes. We find that strong exploitation drive enhances organizational outcomes under low exploitation urgency but also degrades organizational outcomes in some other situations. Moreover, a *deficit* in exploitation capability may be compensated with strong exploitation drive and vice versa under high exploitation urgency. We also illuminate further the role of the strategic context determination process in the transition from exploration to exploitation in strategic innovation projects. Our integrated research findings corroborate the usefulness of the genetic computational modeling approach and suggest areas for further collaboration of computational researchers and field researchers.

Keywords: transition of exploration to exploitation, strategic exploitation challenges, exploitation urgency, exploitation drive, exploitation capability, organizational learning outcomes.

1.0 Introduction

1.1 *Exploration-to-Exploitation Transition Timing in Strategic Projects*

Balancing exploration and exploitation (March, 1991) is a critical challenge of organizational learning and adaptation (Levinthal, 2021). An important aspect of this challenge involves determining the timing of the transition from exploration to exploitation. This would seem to be especially important for *strategic* innovation projects. Such projects involve significant

commitment of resources (Ghemawat, 1991) and the development of products based on competencies that extend, or are significantly different from, the company's current distinctive competencies (Burgelman, 1983; Crossan and Apaydin, 2010; Dominguez-Escrig et al., 2019; Kacperczyk, 2012; Mirabeau and Maguire, 2014). Their success and failure may have serious consequences for the organization.

Deciding on the timing of the transition from exploration to exploitation is a prerogative of the top management. However, the actual decision also depends on the perceptions of differentially positioned executives in a multilevel strategic decision context (e.g., Bower and Gilbert, 2005; Burgelman et al., 2023; Csaszar and Levinthal, 2016; Levinthal, 2021). Biased perceptions may emerge, sometimes inadvertently, because of the dispersed relevant knowledge and commitments that are often enacted simultaneously (Reitzig and Sorenson, 2013). The strategic ambiguity caused by such different perceptions has potentially important consequences for project-level transition from exploration to exploitation, and for organization-level learning.

Previous field research (Burgelman and Aaltonen, 2024) highlighted the managerial challenges related to this ambiguity and informed the research question that motivated the computational study reported in this paper: *How do variations in exploitation urgency, exploitation drive, and exploitation capability affect organizational learning outcomes during the transition from exploration to exploitation in a strategic innovation project?* We examine the conditions under which managers should prioritize exploitation drive (the extent of mobilizing personnel for a strategic project) versus exploitation capability (the extent of information-processing ability for a strategic project) when facing varying degrees of pressure to accelerate the transition from exploration to exploitation (exploitation urgency). We seek to determine how managerial actions necessary to enhance exploitation drive are different from managerial actions necessary to enhance exploitation capability and offer some

guidance for dealing with potential deficits in either exploitation drive or in exploitation capability.

1.2 Field Research Findings of Exploration-to-Exploitation Transition Challenges

To provide an empirical data foundation for our computational analysis we use field research findings of the managerial challenges associated with exploration-exploitation transition timing manifest in the exemplary case (Siggelkow, 2007) of Pharma (disguised name), a multi-business European pharmaceutical company (Burgelman and Aaltonen, 2024). This case study is explained in more detail in Section A of the *Appendix*. It helped document the strategic exploitation challenges Pharma faced in the ambiguous process of transitioning from exploration to exploitation associated with radically new biotechnology-based innovation projects in its therapeutics business unit, and their implications for organizational learning and corporate survival in the rapidly consolidating global pharmaceutical industry.

In general, strategic innovation projects move through a stage-gate process (Cooper, 2008). Progress by clearing the applicable bar in successive stages often drives the escalation of resource commitments (Keil et al, 2000; Klingebiel and Esser, 2020; McNamara et al, 2002; Schmidt and Calantone, 2002; Slesman et al, 2018) and/or impedes project termination (Brockner et al, 1986; Green et al, 2003; Cooper, 2008). The stage-gate process of strategic pharmaceutical projects typically comprises five stages: a preclinical stage (A), followed by four stages (B–E) associated with regulatory-defined I–IV stages (please see *Appendix A* for detailed descriptions of the stages).

The study of Pharma indicated that corporate-level executives from R&D, marketing, finance and operations had significantly different perceptions about the appropriate timing of the transition from exploration to exploitation of strategic projects in this new and relatively unfamiliar product-market category—biotechnology-based innovation in the therapeutics

business unit – than of more familiar innovation projects of Pharma’s core gynecology and diagnostics business units.

Lagging organizational learning about novel biotechnology-related project capabilities—necessary to move from exploration¹ to exploitation (e.g., Brady and Davies, 2004)—had two deleterious outcomes. First, it impeded effective timing of progression of early-stage therapeutic oncology projects prioritized in the stage-gate resource allocation process to later development stages. Second, it impeded the timely termination of failing projects. These developments provide a plausible explanation for why Pharma top management was unable to achieve its strategic intent to grow fast enough to withstand consolidation in the global pharmaceutical industry. The company was acquired (at a hefty premium) by a larger European pharmaceutical company in 2006.

1.3 Computational Analysis of Exploitation Urgency Impact on Exploitation Challenges

Informed by field research findings of our exemplary case study about Pharma’s fading corporate survival prospects, we use computational analysis to examine the *strategic exploitation challenges* associated with transitioning too soon from exploration to exploitation. We examine managerial decisions, highlighted in the study of Pharma, concerning the appropriate levels of *exploitation drive* (extent of mobilization of personnel) and *exploitation capability* (extent of information-processing ability) to deploy in a strategic innovation project. Reaching beyond the findings of the study of Pharma we identify and analyze the importance of the presence of variable extent *exploitation urgency*—related to the speed at which a project is moved out of the exploration stage.

While there is broad agreement in extant research that exploitation is favored over exploration (e.g., Greve, 2007; He and Wong, 2004; Sirén et al., 2012), there appears to be

¹ A move from the pre-clinical stage to later stages is considered as a move from exploration to exploitation, i.e., a move from “research” to “development”.

little guidance from research elaborating on *why* increasing exploration—instead of strongly favoring exploitation—improves longer-term outcomes, and what recourses apply when increasing the quantum of exploration is not feasible. In other words, how top management can achieve exploration-exploitation balance (e.g., Lavie et al., 2010; Levitt and March, 1988; Levinthal and March, 1993; March 1991; Schmitt et al., 2010) is an open question.

In what follows we attempt to find some answers to help fill these gaps in knowledge. We integrate computational research using modest modifications of March's (1991) genetic model with prior field research findings to quantitatively determine organizational learning outcomes associated with the timing of transitioning from exploration to exploitation in a strategic innovation project under stylized conditions. We find that, under conditions of low exploitation urgency, strong exploitation drive obtains superior organizational learning outcomes compared to outcomes from weak exploitation drive. However, under conditions of high exploitation urgency, strong exploitation drive in conjunction with strong exploitation capability obtains the most inferior organizational learning outcomes. Further, under conditions of high exploitation urgency, organizational learning outcomes are minimally penalized by exploitation drive deficit if exploitation capability is strong; or by exploitation capability deficit if exploitation drive is strong.

In summary, our paper makes the following contributions. First, we integrate two long-standing theoretical perspectives in organization theory and strategic management: (i) computational-based theory of exploration and exploitation in organizational learning and (ii) field research-based managerial agency theory of the role of strategic innovation in organizational strategy. Second, with modest extensions of March's (1991) model of exploration and exploitation we disentangle the effects on organizational learning outcomes of low/high exploitation urgency in relation with conjunctions of weak/strong exploitation drive and weak/strong exploitation capability. Third, our findings also illuminate further the

danger of top management proactively but prematurely determining the incipient strategic context of a strategic innovation project and signaling their exploitation urgency, which curtails exploration and learning of new necessary competencies. These contributions support the view that combining field-based research and computational-based research can help identify new areas for research that potentially extend the reach of both.

2.0 Theoretical parameters elicited from the exemplary case

The story of Pharma leads us to consider certain nuances in the exploration-exploitation phenomenon. First, the stage-gate nature of the phenomenon suggests that the context is one where a project initiates with emphasis on exploration (e.g., the pre-clinical phase).

Exploitation dominates at a later point in time. The timing of exploration-exploitation transition is an important parameter because too little or too much of either exploration or exploitation is unlikely to be beneficial to a company (March, 1991). Observing pharma top-management's eagerness to get therapeutic projects into high-visibility development (exploitation) stage as early as possible, we introduce the term 'exploitation urgency'. Higher the *exploitation urgency* sooner is a project transitioned from exploration to exploitation.

The Pharma study also evidences high degree of mobilization of personnel and resources for projects of the favored-for-growth therapeutics business unit. We cite the construct concerning mobilization as 'exploitation drive'. We reason that strong *exploitation drive* is likely to have beneficial outcomes if the underlying product obtained sufficient degree of exploration (research), i.e. if exploitation urgency was low.

We designate the information processing capabilities of mobilized personnel deployed in the exploitation (clinical) phase as 'exploitation capability'. We further reason that, deployment of strong exploitation drive and strong exploitation capability to a product that received insufficient exploration—applicable under conditions of high exploration urgency—is quite likely to be a recipe for failure. Further, the possibility that a product received

inadequate exploration is quite real, given that prioritization of high number of early-stage projects is likely to constrain the extent of exploration (research) an individual project is likely to receive.

The considerations discussed above lead us to appreciate certain strategic exploitation challenges. These concern the extent of exploitation drive and exploitation capability to deploy given varying extent of exploitation urgency, in a stage-gate phenomenon. We take up formally describing the theoretical concepts guiding our computational analysis next².

2.1 Exploration-Exploitation Transition and Exploitation Urgency

The timing of the transition from exploration to exploitation associated with strategic innovation projects may be affected by external and/or internal forces. These forces may determine *exploitation urgency* as signaled by top management and perceived as such—correctly or incorrectly—by middle/senior management. With high exploitation urgency an innovation project is given a shorter time for exploration. High exploitation urgency may be noticeable from top management strategic intent of growing fast and signaling the same to the external world by pushing projects into the high-visibility exploitation stages.

Exploitation urgency is viewed in our study as a distinct force that is inversely related to the time allotted to exploration in a strategic innovation project. Exploitation urgency is *high* if top management supports moving rapidly from exploration (research) to exploitation (development). Allowing exploration to continue for an extended duration connotes *low* exploitation urgency³.

2.2 Exploitation Drive and Exploitation Capability

² As a matter of theoretical refinement, we have split the conjugation of high mobilization and high urgency to move a project from exploration to exploitation designated as ‘exploitation drive’ in Burgelman and Aaltonen (2024) into separate constructs ‘exploitation drive’ and ‘exploitation urgency’. Likewise, the conjugation of information-processing ability and organizational learning cited as ‘exploitation capability’ are separated into distinct constructs, ‘exploitation capability’ and ‘organizational learning’.

³ In the context of the *Pharma* case, the additional time in exploration could also be utilized for developing better grasp of the new medical-related capabilities applicable to oncology and cardio-vascular therapeutics, such that management avoided the pitfall of inappropriate allocation of the same on account of inexperience.

Exploitation drive refers to the extent of mobilization of personnel associated with a strategic innovation project at any given point in time during the progress of the project. Exploitation drive is *strong* if there is mobilization of higher proportion of team members for exploitation activities. Exploitation drive is *weak* otherwise.

Exploitation capability refers to the extent of information-processing ability—associated with learning from repositories of knowledge (the ‘organizational code’)—mustered by the innovation team members mobilized for exploitation tasks. If exploitation capability is weak, it may in certain circumstances reduce the negative effects of groupthink because less of the existing organizational knowledge—potentially deficient in relation to the innovation project—is applied; but under other circumstances, strong exploitation capability may be beneficial for organizational learning outcomes⁴.

2.3 *Exploitation Drive Deficit and Exploitation Capability Deficit*

In general, *exploitation drive deficit* can materialize on account of (a) difficulty in mobilization of increased numbers of personnel for exploitation tasks, (b) simple neglect given that other projects are in top management focus, and (c) the nature of a given project itself, whereby mobilization of higher numbers of personnel is ineffective or counterproductive. The outcome of *exploitation drive deficit* is nominally measured by subtracting the level of organizational learning (or code knowledge, elaborated below) obtainable by weak exploitation drive from that obtainable by strong exploitation drive.

Exploitation capability deficit occurs when a strategic innovation project experiences limitations of information-processing abilities of personnel deployed for exploitation tasks.

For instance, when the organizational personnel have too many other preoccupations it

⁴ In the *Pharma* case, refinement of the learning regarding conducting complicated clinical trials in oncology and cardio-vascular therapeutics would manifest as positive effect of strong exploitation capability. Weak exploitation capability would connote slower appreciation of the idiosyncrasies that make clinical trials in oncology and cardio-vascular therapeutics different from other clinical trials familiar to company personnel.

becomes difficult to devote time to refine the new capabilities. In our study, the outcome of *exploitation capability deficit* is nominally measured by subtracting the level of organizational learning obtainable under weak exploitation capability from that obtainable by strong exploitation capability.

2.4 Strategic Context Determination and Organizational Learning Outcomes

Prior research (Burgelman, 1988) has noted that *strategic context determination* and *organizational learning* are intrinsically related. There is separation of action and cognition over time. Managers learn by doing (e.g., they do more of the same when successful); sense-making (as to what may have contributed to success) is carried out in defining the cognition constructs and noting them in strategy documents. In other words, the “experience base” contributes to development of a strategy framework. Thereby, (correct or incorrect) organizational learning could be a driving force behind much intra- and inter-organizational change (Holmqvist, 2004).

To decide whether to sustain support for a strategic innovation project, top management can use the strategic context determination process and *modify the rate of exploitation* to assess the significance of changes in organizational learning in the innovation project (Burgelman and Chanda, 2024). In the present study, we examine how strategic context determination for a strategic innovation project secures augmentation of organizational learning outcomes by the interplay of exploitation drive and exploitation capability, in a backdrop of varying exploitation urgency.

In our study the *organizational learning outcome* is measured by the extent of match with the knowledge required to attain the innovation project target. In a strategic innovation project, a configuration of stylized parameters that obtains higher organizational learning (i.e., greater match between the values in the organizational code dimensions with the values in the corresponding dimensions of the innovation target) is considered superior to a

configuration that attains a lower level of organizational learning. Effectively, organizational learning is a measure of code *knowledge* residing in the repository of project level knowledge. Code knowledge or organizational learning—also referred to as *organizational knowledge* in the literature—comprises information stored in databases, user manuals, rules, forms, standard operating procedures, past and current strategic plans (Burgelman and Chanda, 2024; March, 1991). A project team member mobilized for exploitation tasks learns from the organizational code/strategy (at a certain rate, given by the rate of exploitation), subject to limits of information-processing ability which as mentioned above we designate as *exploitation capability*. The organizational code learns from the knowledgeable members, and the probability of the learning being committed to is higher if the majority vote of highly knowledgeable team members is higher. We assume that the organizational code thereby encodes recipes that are deemed to be successful; some incorrect knowledge, however, will seep into the organizational code, simply because the majority view is not infallible.

3.0 Computational model

The simulation model in this study draws from the conceptual specifications of the genetic algorithm model presented by March (1991) as elaborated by Chanda and Miller (2019)⁵. The simulation model is closely similar to the model presented in Burgelman and Chanda (2024)⁶. The constructs of *exploitation urgency*, *exploitation drive* and *exploitation capability* are inspired by recent field research studying the fate of strategic innovation projects in a European pharmaceutical company, over an extended period (Burgelman and Aaltonen, 2024).

⁵ Important to note, we use extensions of March's (1991) genetic computational model rather than the *n-arm-bandit* model or Kauffman's (1993) *NK* model, which are predominantly single-agent models that are less suitable for studying the distinct processes and outcomes at both the individual and organizational levels that we focus on here. To clarify the reasons for this choice we provide a comparison of the *NK*, Bandit and the March (1991) genetic model in Section **E** of the *Appendix*.

⁶ In Section **B** of the *Appendix*, the evolution from March's (1991) model to Burgelman and Chanda (2024) to that used in the current manuscript is discussed in detail.

The strategic innovation project team (called innovation team in most places in the remainder of the paper) comprises N members. Each member has M belief dimensions. The innovation team assimilates knowledge ('code knowledge') in the form of *organizational learning* in an entity, the organizational code (**OC**), which also has M dimensions⁷. Initially **OC** is populated with all values "0", signaling neutral values or "no opinion". The external reality (**R**) comprising the innovation target is a string with M bits. At the beginning, **R** is populated with values "-1" and "+1", where each value has 50% probability of materialization. Each simulation experiment runs for T time-steps. In each time-step the level of code knowledge (i.e. **OC**'s knowledge) is computed as the number of dimensions having matching values in **OC** and **R**, divided by M , the number of dimensions.

At the beginning of simulation experiments, the *innovation team knowledge* or *organizational member knowledge*⁸ (Burgelman and Chanda, 2024) comprising the latent knowledge in the heads of the innovation team members is vastly misaligned with the knowledge required to attain the innovation target. This is fashioned by first populating the member belief dimensions with values "-1", "0" and "+1" with equal (one-third) probability (March, 1991), and then overwriting (a randomly chosen) fifteen percent of the dimensions with values opposite to that in the corresponding dimension of **R** (Burgelman and Chanda, 2024; Chanda et al., 2018)⁹.

The knowledge carried by an individual member is computed as the number of dimensions for which a value in a member-string-cell matches the corresponding value in **R**,

⁷ Burgelman and Chanda (2024) refer to this placeholder as the "commandeered section of the official organizational code" (**CSOC**) to sharply distinguish the organizational knowledge developed by an autonomous innovation project team from other organizational knowledge in the broader organization. In this paper, given that top management supports the radical innovation effort, we refer to the **CSOC** as **OC**.

⁸ Organizational member knowledge is alternately referred to as "collective human capital" at other places, e.g., von Nordenflycht (2011), Chanda et al. (2018) and Chanda (2024).

⁹ Effectively, the initial misalignment between the team knowledge and the innovation target is implemented by having the ratio of the initial average level of incorrect knowledge to correct knowledge in the team increase from 1:1 to 1.5:1.

We signify *exploitation capability* to denote the extent of information-processing ability possessed by any member mobilized for carrying out exploitation tasks. Here, information-processing capability is a function of the number of dimensions processed by a member (undertaking exploitation activity) and the number of time-steps for which such processing is carried out, in the course of the strategic innovation project. Thus, *exploitation capability* is higher if a greater fraction of knowledge dimensions is considered for update from **OC** and/or if such consideration is carried out in a greater fraction of time-steps¹¹.

The third flow concerns the rate of inflow of heterogeneous knowledge from outside the organization. It is mapped to the parameter $p3_alt$. A non-zero value of this parameter maps to the rate of exploration, by members of the strategic innovation project. A non-zero value (say ten percent) signifies the probability that belief dimensions of a randomly-chosen subset (one-fourth) of organizational members are randomized as a result of acquiring heterogeneous new knowledge from outside the organization. *Exploitation urgency* is low if exploration is kept turned ON (by having a non-zero value in $p3_alt$) for a major duration of the simulation run; if exploration is withdrawn early, *exploitation urgency* is high.

The outcome variable is *organizational learning* or *code knowledge*. *Organizational learning* (*code knowledge*) is computed by dividing the number of matches between corresponding values in **R** and **OC** by the number of dimensions (M). A configuration of parameter values that produces higher organizational learning (code knowledge) is considered superior to any other configuration that attains lower organizational learning (code knowledge). The reported results are averages over 10,000 runs of a given experiment, with stochastically varying inputs for the parameters $p1$, $p2$, $p3_alt$, and for the probabilities

¹¹ In the simulation experiments, we use equal values for fractions signifying scope (fraction of knowledge dimensions processed) and frequency (fraction of time-steps for which processing is carried out) underlying the exploitation rate ($p1$). For example, we use the square-root of 0.10 as values for frequency and scope to signify low (10%) information processing capability; to signify moderate capability (50%) we use square-root of 0.50; we use square-root of 0.90 to signify high exploitation capability (90%).

We observe that higher exploitation drive has a positive impact on outcomes of organizational learning (code knowledge) when exploitation capability is weak (the curve “Weak Capability”). This happens because increasing the proportion of innovation team members mobilized for exploitation (strong exploitation drive) in conjunction with lower fraction of knowledge dimensions considered for update from the organizational code **OC** and/or in a lower fraction of time-steps (weak exploitation capability) maintains higher degree of diversity of member knowledge on account of slower rate of socialization; this has a positive impact on organizational learning outcomes. In this case, our results also show that a *severe* deficit in exploitation capability can be compensated through strong exploitation drive.

We further observe that exploitation drive has a negative impact on organizational outcomes for moderate and strong levels of exploitation capability (the curves “Moderate Capability” and “Strong Capability”). This happens because heterogeneity of knowledge in the innovation team personnel is driven out rapidly through stronger socialization (members rapidly learning the knowledge residing in the organizational code), leading to overall lower organizational learning. In other words, while a low extent of mobilization of organizational personnel for exploitation tasks (i.e., weak exploitation drive) enables longer preservation of valuable diverse knowledge of members—enabling attainment of higher level of organizational knowledge—under strong exploitation drive in conjunction with moderate or high exploitation capability, the members become socialized quicker, leading to faster erosion of diversity of knowledge, and this results in attaining a lower level of organizational learning (code knowledge).

In **Figure 3** we present variation of *organizational learning (code knowledge)* with *exploitation drive* and *exploitation capability*, in presence of *exploration* involving intake of

heterogeneous knowledge from outside the company, throughout the life time of a strategic innovation project.

>>>>>>>> Insert **Figure 3** about here <<<<<<<<<<<<<<

We observe that increasing exploitation drive in conjunction with stronger exploitation capability helps attain superior organizational learning outcomes. However, the effect of stronger exploitation drive is somewhat different from that observed in Figure 2. We observe that stronger exploitation drive leads to better outcomes initially, and flattens thereafter. This happens because there are two sources of valuable heterogeneous knowledge in this case (a) team members' latent knowledge and (b) inflow from outside sources. Higher mobilization for exploitation tasks assists in greater assimilation of knowledge. Organizational learning increases till the point of time that the erosion of variety of member knowledge is more than made up by the inflow of variety from outside.

4.2 *Computational analysis with low & high exploitation urgency: Formal propositions*

In **Figure 4** we present organizational outcomes under low and high *exploitation urgency* for strategic innovation projects having combinations of strong and weak *exploitation capability* and strong and weak *exploitation drive*. Under *low exploitation urgency*, exploration is allowed for an extended duration (80 out of 100 time-steps); under *high exploitation urgency*, exploration is given a far shorter run (20 out of 100 time-steps).

>>>>>>>> Insert **Figure 4** about here <<<<<<<<<<<<<<

The left-hand panel of Figure 4 shows that under conditions of *low exploitation urgency*, a conjunction of *strong exploitation capability* and *strong exploitation drive* obtains the most superior organizational learning outcomes. Figure 4 also shows the ranking of the computational results of organizational learning outcomes for low exploitation urgency: Strong Capability/Strong Drive, followed by Weak Capability/Strong Drive, then by Strong Capability/Weak Drive and finally Weak Capability/Weak Drive. Thus, we have:

Proposition 1. *Under conditions of low exploitation urgency, strong exploitation drive obtains superior organizational learning outcomes compared to outcomes from weak exploitation drive.*

The right-hand panel of Figure 4, in contrast, shows that under conditions of *high exploitation urgency* a conjunction of *strong exploitation capability* and *strong exploitation drive* obtains markedly inferior outcomes. Figure 4 also shows the ranking of the computational results of organizational learning outcomes for high exploitation urgency: Weak Capability/Strong Drive, followed by Strong Capability/Weak Drive, then by Weak Capability/Weak Drive and finally Strong Capability/Strong Drive. Thus, we have:

Proposition 2. *Under conditions of high exploitation urgency, strong exploitation drive in conjunction with strong exploitation capability obtains the most inferior organizational learning outcomes.*

The information presented in Figure 4 can be recast to study organizational outcomes under conditions of exploitation capability deficit and exploitation drive deficit. A penalty arising from a ‘deficit’ in a parameter (capability or drive) is calculated by subtracting the extent of organizational learning (code knowledge) attained under low values of the parameter from the organizational learning (code knowledge) attainable under high values of the parameter. If the result is a positive number (i.e., if a positive value appears in the vertical axis in Figure 5 or Figure 6), we infer that there is a ‘penalty’ attributable to the ‘deficit’. Otherwise, if the result is a negative number, there is no penalty; in fact, the situation is beneficial to the organization.

>>>>>>>> Insert **Figure 5** about here <<<<<<<<<<<<<<

Figure 5 shows the extent of *penalty* accruing from *exploitation capability deficit*. The values on the vertical axis are obtained by subtracting the organizational learning (code knowledge) attained under *weak exploitation capability* from that attained under *strong exploitation capability*.

>>>>>>>> Insert **Figure 6** about here <<<<<<<<<<<<<<

Figure 6 shows the extent of *penalty* accrued from *exploitation drive deficit*. The values on the vertical axis are obtained by subtracting the organizational knowledge attained under *strong exploitation drive* from the organizational knowledge attained under *weak exploitation drive*.

Under conditions of high exploitation urgency, we observe (in **Figure 5**) that, deficit in exploitation capability does not yield performance penalty when in conjunction with a strong exploitation drive; moreover, we observe (in **Figure 6**) that deficit in exploitation drive does not yield performance penalty when in conjunction with a strong exploitation capability. In both cases, performance is superior, rather than inferior. Thus, we suggest:

Proposition 3. *Under conditions of high exploitation urgency, organizational learning outcomes are not penalized by exploitation drive deficit if exploitation capability is strong; or by exploitation capability deficit if exploitation drive is strong.*

Effectively, enhancing either exploitation drive or exploitation capability can make up for deficit in the other. These act as substitutes, under conditions of high exploitation urgency, with the proviso that simultaneously strong levels of both are undesirable¹⁴.

4.3 Boundaries (Limitations) of the study

First, we develop propositions about organizational learning outcomes attained by an *individual* strategic innovation project under stylized conditions of operations. In real-world firms, often a number of such projects are simultaneously in progress. Hence, this study refrains from furnishing inferences about the fate of a *portfolio* of innovation projects (Si et al., 2022).

Second, we assume that organizational personnel not mobilized for exploitation tasks in a given time-step do not contribute to membership of the team of knowledgeable members that advise the organizational code, in that time-step. In other words, members that are

¹⁴ Detailed, model-based explanations of the key results are provided in Section **C** of the *Appendix*. Additional discussion focusing on outcomes under extreme values of the p_2 parameter appear in Section **D** of the *Appendix*.

inactive from the standpoint of exploitation tasks are excluded from participating in update of organizational learning (code knowledge). This arrangement enables the code to obtain high quality inputs from members that are currently active. This leads to attainment of a higher extent of organizational learning (code knowledge). In reality though, even inactive members may demand that their voices be paid attention to, in knowledge codification.

Third, our study concerns mapping the transformation of *latent* knowledge (Hargaddon and Fanelli, 2002) of individuals—arising from prior life experiences or through inflow of heterogeneous knowledge from outside the company— to codification into organizational learning (code knowledge). However, it leaves out the handling of *tacit knowledge* (Miller et al., 2006)—a kind of knowledge that resists codification and is passed on through apprenticeship, learning-by-doing and other means.

5.0 Discussion and Implications

5.1 *Computational results inform strategic exploitation challenges*

Our computational findings indicate that exploitation drive has a big impact on augmenting organizational learning outcomes, but that the direction of the impact (positive or negative) is greatly dependent on both exploitation urgency and exploitation capability. In view of Proposition 1 (please also refer to the left-hand panel of Figure 4), under conditions of low exploitation urgency, an innovation team with strong exploitation capability in conjunction with strong exploitation drive obtains superior organizational learning outcomes. It thereby displays “exploitation balance” (Burgelman and Aaltonen, 2024). If a team’s exploitation capability is weak, a conjunction with strong exploitation drive produces outcomes that are slightly inferior, signifying a “capability deficit” (Burgelman and Aaltonen, 2024). A team’s strong exploitation capability in conjunction with weak exploitation drive produces outcomes somewhat more inferior. It faces a “drive deficit” (Burgelman and Aaltonen 2024). If both

exploitation capability and exploitation drive are weak, the outcomes are most inferior. The team experiences an “exploitation trap” (Burgelman and Aaltonen, 2024).

In view of Proposition 2 (please also refer to the right-hand panel of Figure 4), under conditions of high exploitation urgency, “exploitation balance” is unattainable. A strategic innovation team deploying strong exploitation capability and strong exploitation drive obtains the most inferior organizational learning outcomes. In this situation there is the least time to learn by experimentation and add to the organizational code/strategy the new distinctive competencies necessary for success. Even an ‘exploitation trap’—weak exploitation capability in conjunction with weak exploitation drive—obtains better outcomes. Superior outcomes are obtained under a combination of *either* strong exploitation drive and weak exploitation capability *or* weak exploitation drive and strong exploitation capability. This signifies that deficiency in exploitation capability can be compensated by strong exploitation drive, and vice versa.

Propositions 1 and 2 indicate that it probably would have been better for the European pharmaceutical company to resist the impulse to disproportionately prioritize many early-stage oncology therapeutics projects in the resource allocation process because this decision reduced the extent of exploration available to an individual project. Continued exploration of the necessary new biotechnological competencies necessary to move projects successfully to later development stages would probably have produced superior organizational learning outcomes, leveraging the company’s exploitation drive and exploitation capability.

Proposition 3 formalizes the finding that under conditions of high exploitation urgency, exploitation capability deficit can be compensated by strong exploitation drive, and exploitation drive deficit can be compensated by strong exploitation capability. This novel insight augments received theory about bounded rationality, such that managers, already

diversity of knowledge of the innovation team members can continue to be augmented through exploration. In conjunction with weak (ten percent) information-processing capability (learning from the existing organizational knowledge) this still produces strong organizational learning outcomes, taking advantage of the relatively longer persistence of knowledge heterogeneity in the team members. However, in conjunction with strong (ninety percent) information-processing capability this has a large negative impact on the magnitude of organizational learning outcomes. This latter conjunction indicates the danger of exploitation urgency, which, in the face of lacking relevant new necessary knowledge, causes the deployment of existing irrelevant organizational knowledge to strategic innovation projects.

Our findings indicate the need for further research to examine the forces that determine exploitation urgency, especially in relation to organizational adaptation and longevity. In the case of the European pharmaceutical company, for instance, capital market forces, changing ecosystem forces, and product-market forces affected the different managerial levels (respectively higher, middle and lower) differently, but all contributed to top management proactively but prematurely determining the incipient strategic context of the new oncology therapeutics projects and signaling their exploitation urgency. The unexpected negative outcomes of the exploitation urgency-determined strategic resource allocation process—which curtailed exploration and associated learning of new necessary medical-related capabilities by the oncology therapeutics projects—caused attenuation and fizzling out of their incipient strategic context, and stymied top management strategic intent to rapidly grow the new business and maintain the company’s independence (Burgelman and Aaltonen, 2024).

5.3 *Implications for received theory and further research*

Investing in exploitation at the cost of exploration is often tempting for a firm because the outcomes of exploitation are proximate and certain (March, 1991). Greve (2007) finds that exploitation of current capabilities often reduces exploration of new capabilities. Sirén et al. (2012) argue that learning processes favoring assimilation of exploitative knowledge for commercial ends have been found to form exploitation traps, which are also referred to as competency traps (Levinthal and March, 1993; Wu and Shanley, 2009). In this type of exploitation trap high levels of exploitation dominate the strategic learning capability of a firm, restricting the explorative innovations (He and Wong, 2004). Not providing a sufficient time horizon for allowing a strategic innovation project to demonstrate its potential for meaningfully augmenting organizational knowledge constitutes an important hazard of top management impatience (Burgelman and Chanda, 2024).

Notwithstanding the broad agreement that exploitation is favored, little extant research proposes mechanisms explaining *why* increasing exploration—instead of strongly favoring exploitation—improves longer-term outcomes. The insights into strategic exploitation challenges facing top management generated by our computational study help filling this gap.

While our computational analysis has generated novel insights beyond those offered by prior field research of strategic exploitation challenges, it also suggests that further field-based studies—aimed at increasing our understanding of how top management can achieve exploration-exploitation balance—would be vitally important. Such studies should seek out additional “exemplary cases” (Siggelkow, 2007) to investigate what actions top management can take in order to re-shape the strategic innovation project narrative favorably, such that projects are given extended time in exploration even though they are likely to have lower visibility to external stakeholders, such as market analysts, that top management wishes to impress to gain continued support. Such field-based studies could also begin to examine the

relationship between organizational learning outcomes and organizational adaptation, knowledge about which is currently scarce. Further computational-based research could model the insights gained from these field-based studies and simulate their impact on organizational learning outcomes and adaptation, thereby generating richer and deeper insight and enhancing further the productive collaboration between practitioners of these different research methods for further advancing strategic management knowledge.

6.0 References

- Bontis N, Crossan MM, Hulland J (2002) Managing an organizational learning system by aligning stocks and flows. *Journal of Management Studies* 39(4):437–469.
- Bower JL, Gilbert CG (2005) *From Resource Allocation to Strategy*. (Oxford University Press, New York).
- Brady T, Davies A (2004) Building project capabilities: from exploratory to exploitative learning. *Organization Studies* 25(9):1601–1621.
- Brockner J, Houser R, Birnbaum G, Lloyd K, Deitcher J, Nathanson S, Rubin J (1986) Escalation of commitment to an ineffective course of action: The effect of feedback having negative implications for self-identity. *Administrative Science Quarterly*, 31, 109–126.
- Burgelman RA (1983) Corporate entrepreneurship and strategic management: Insights from an process study, *Management Science* 29(12):1349–1364.
- Burgelman RA (1988) Strategy making as a social learning process: The case of internal corporate venturing. *Interfaces* 18(3):74–85.
- Burgelman RA (2002) Strategy as vector and the inertia of co-evolutionary lock-in. *Administrative Science Quarterly* 47(2):325–357.
- Burgelman RA, Snihur Y, Thomas LDW (2023) *Strategy-Making and Organizational Evolution: A Managerial Agency Perspective* (Cambridge University Press, Cambridge).
- Burgelman RA, Aaltonen P (2024) Fading corporate survival prospects: Impact of co-selection bias in resource allocation on strategic intent. Working Paper 4191. *Stanford Graduate School of Business*. Forthcoming in *Strategic Management Journal*.
- Burgelman RA, Chanda SS (2024) Autonomous strategic behavior, organizational learning and top management support: Re-examining field research with computational modeling. Working Paper 4176. *Stanford Graduate School of Business*. Forthcoming in *Strategic Management Review*.
- Chanda SS (2024) Enabling creative destruction: A view from the perspective of managerial control. *Journal of Contemporary Business Research*. DOI: 10.1177/3049513X241233064.
- Chanda SS, Miller KD (2019) Replicating agent-based models: Revisiting March’s exploration-exploitation study. *Strategic Organization* 17(4):425–449.
- Chanda SS, Ray S, McKelvey B (2018) The continuum conception of exploration and exploitation: An update to March’s theory. *M@n@gement* 21(3):1050–1079.

- Cooper RG (2008) The stage-gate idea-to-launch process – update, what’s new and NexGen systems. *Journal of Product Innovation Management* 25, 213–232.
- Crossan MM, Apaydin M (2010) A multi-dimensional framework of organizational innovation: A systematic review of the literature. *Journal of Management Studies* 47(6): 1154–1191.
- Csaszar FA, Levinthal DA (2016) Mental representation and the discovery of new strategies. *Strategic Management Journal* 37(10): 2031–2049.
- Dierickx I, Cool K (1989) Asset accumulation and sustainability of competitive advantage. *Management Science* 35(12):1504–1511.
- Dominguez-Escrig E, Mallén-Broch F F, Lapiedra-Alcami R, Chiva-Gomez R (2019) The influence of leaders’ stewardship behavior on innovation success: The mediating effect of radical innovation. *Journal of Business Ethics* 159(3):849–862.
- Ghemawat P (1991) *Commitment: The Dynamic of Strategy*. Free Press.
- Green SG, Welsh MA, Dehler G (2003) Advocacy, performance, and threshold influences on decisions to terminate new product development. *Academy of Management Journal* 46, 419–434.
- Greve H (2007) Exploration and exploitation in product innovation. *Industrial and Corporate Change* 16(5):945–975.
- Hargadon A, Fanelli A (2002) Action and possibility: Reconciling dual perspectives of knowledge in organizations. *Organization Science* 13(3):290–302.
- He Z-L, Wong P-K (2004) Exploration vs. exploitation: An empirical test of the ambidexterity hypothesis. *Organization Science* 15(4):481–494
- Holmqvist M (2004) Experiential learning processes of exploitation and exploration within and between organizations: an empirical study of product development. *Organization Science* 15(1):70–81.
- Kacperczyk AJ (2012) Opportunity structures in established firms: Entrepreneurship versus intrapreneurship in mutual funds. *Administrative Science Quarterly* 57(3):484–521.
- Kauffman SA (1993) *The Origins of Order: Self-organization and Selection in Evolution*. Oxford, UK: Oxford University Press.
- Keil M, Mann J, Rai A (2000) Why software projects escalate: An empirical analysis and test of four theoretical models. *MIS Quarterly* 24, 631–664.
- Klingebiel R, Esser P (2020) Stage-Gate Escalation. *Strategy Science* 5(4), 311–329.
- Lavie D, Stettner U, Tushman ML (2010) Exploration and exploitation within and across organizations. *Academy of Management Annals* 4(1):109–155.

- Levinthal DA (2021) *Evolutionary processes and organizational adaptation: A Mendelian perspective of strategic management*. (Oxford University Press, New York).
- Levinthal DA, March JG (1993) The myopia of learning. *Strategic Management Journal* 14(S2):95–112.
- Levitt B, March JG (1988) Organizational learning. *Annual Review of Sociology* 14, 319–340.
- March JG (1991) Exploration and exploitation in organizational learning. *Organization Science* 2(1):71–87.
- McNamara G, Moon H, Bromiley P (2002) Banking on commitment: Intended and unintended consequences of an organization's attempt to attenuate escalation of commitment. *Academy of Management Journal* 45(2):443–452.
- Miles RE, Snow CC (1978) *Organization Strategy, Structure and Process* (Mc-Graw-Hill, New York).
- Miller KD, Zhao M, Calantone RJ (2006) Adding interpersonal learning and tacit knowledge to March's exploration-exploitation model. *Academy of Management Journal* 49(4):709–722.
- Mirabeau L, Maguire S (2014) From autonomous strategic behavior to emergent strategy. *Strategic Management Journal* 35(8):1202–1229.
- Reitzig M, Sorenson O (2013) Biases in the selection stage of bottom-up strategy formulation. *Strategic Management Journal* 34, 782–799.
- Schmidt JB, Calantone RJ (2002) Escalation of commitment during new product development. *Journal of the Academy of Marketing Science* 30, 103–118.
- Schmitt A, Probst G, Tushman ML (2010) M@n@gement in times of economic crisis: Insights into organizational ambidexterity. *M@n@gement* 13(3):128–150.
- Si H, Kavadias S, Loch C (2022) Managing innovation portfolios: From project selection to portfolio design. *Production and Operations Management* 31(12):4572–4588.
- Siggelkow N (2007) Persuasion with case studies. *Academy of Management Journal* 50(1):20–24.
- Sirén C, Kohtamäki M, Kuckertz A (2012) Exploration and exploitation strategies, profit performance, and the mediating role of strategic learning: Escaping the exploitation trap. *Strategic Entrepreneurship Journal* 6(1):18–41.
- Sleesman DJ, Lennard AC, McNamara G, Conlon D (2018) Putting escalation of commitment in context: A multilevel review and analysis. *Academy of Management Annals* 12, 178–207.

- Srikanth K, Ungureanu T (2025) Organizational adaptation in dynamic environments: Disentangling the effects of how much to explore versus where to explore. *Strategic Management Journal* 46(1):19–48.
- 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.
- Wu J, Shanley MT (2009) Knowledge stock, exploration, and innovation: research on the United States electro-medical device industry. *Journal of Business Research* 62(4):474–483.

7.0 Tables and Figures

Table 1. Parameters in the simulation model and for robustness checks.

Parameter	Description	Value in experiments	Robustness check
M	The number of bits of (i) reality string, (ii) organizational code string and (iii) member belief string	30	25, 35
N	The number of members in the <i>innovation team</i> attempting radical innovation	50	40, 60
p_1	Member learning rate, or overall rate of exploitation	0.50	0.40, 0.60
p_2	The rate of learning by the organizational code	0.50	0.40, 0.60
p_{3_alt}	The rate of exploration: probability that belief dimensions are randomized, for a subset (25%) of organizational members, when accessing heterogeneous knowledge from outside the organization	0.10	0.08, 0.12
T	The number of time-steps that a given simulation experiment is carried out for	100	90, 110
<i>Exploitation Drive</i>	Proportion of the innovation team members mobilized for exploitation in any given time-step (Weak/ Moderate / Strong)	0.10, to 0.90, in steps of 0.10 (0.10, 0.50, 0.90)	N/A
<i>Exploitation Capability</i>	Product of fraction of knowledge dimensions processed and fraction of time-steps for which such processing is done by a member mobilized for exploitation tasks. (Weak/ Moderate / Strong)	0.10, 0.50, 0.90	N/A
<i>Exploitation Urgency</i>	The point in time (time-step number) at which exploration is stopped, high if earlier, low otherwise	20, 80	[15, 25], [75, 85]
<i>Team member knowledge or Collective Human Capital (CHC)</i>	For the members of an innovation team attempting a radical innovation <i>CHC</i> is low (<i>sub-Marchian CHC</i>); in other words, initial team knowledge is vastly misaligned to the knowledge required to achieve the innovation target. Certain percentage (<i>v. def</i> = 15%) of member beliefs—originally created by random draw from $\{-1, 0, +1\}$ —are overwritten with values opposite that of values in corresponding bit positions of the initial reality string, at the beginning of a simulation experiment.	0.15	0.14, 0.16

Table 2. Strategic context determination and organizational learning outcomes

		Innovation Team Exploitation Features	
		<i>Strong Drive & Weak Capability</i>	<i>Strong Drive & Strong Capability</i>
Top Management Urgency	<i>Low</i>	0.79 organizational learning	0.82 organizational learning
	<i>High</i>	0.77 organizational learning	0.5 organizational learning

Figure 1. Schematic diagram for stocks and flows in the model

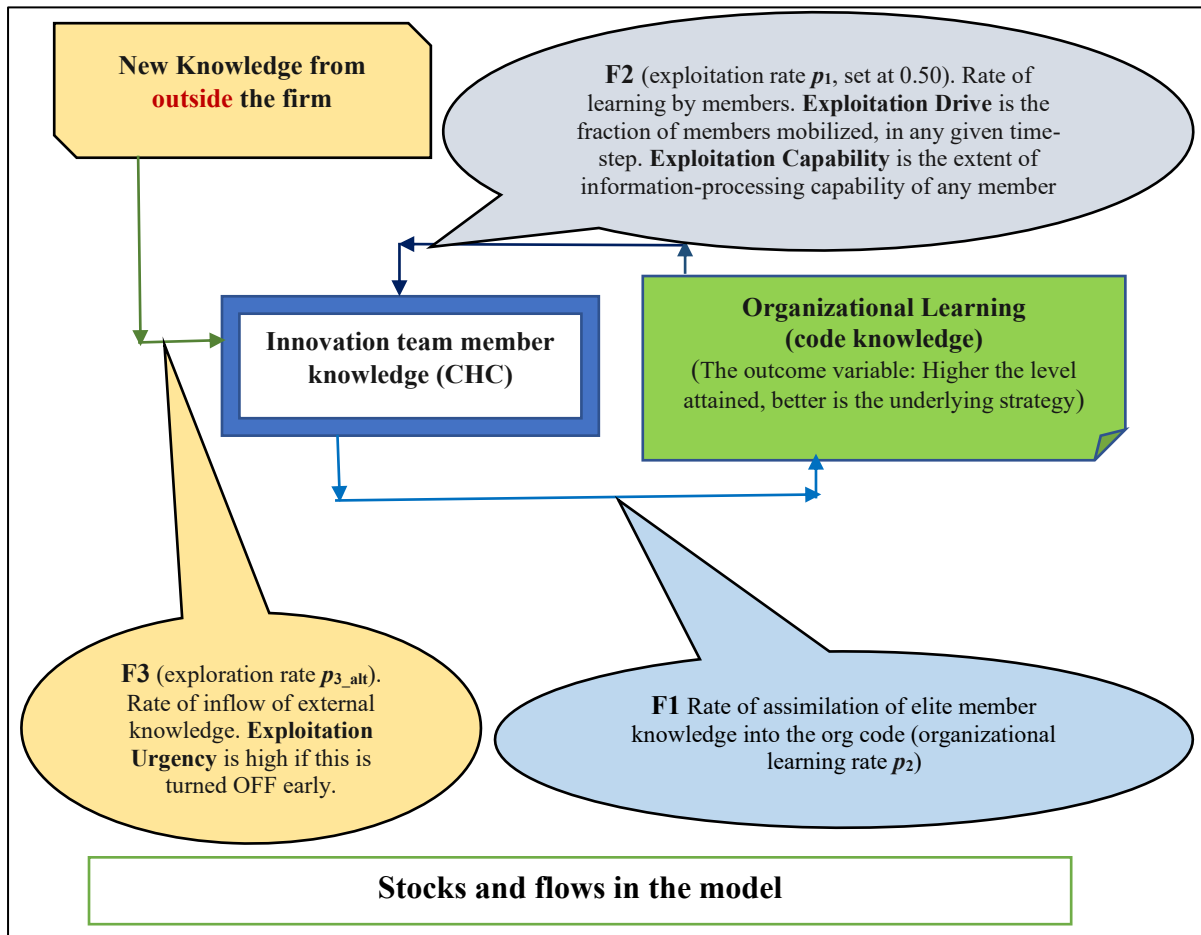
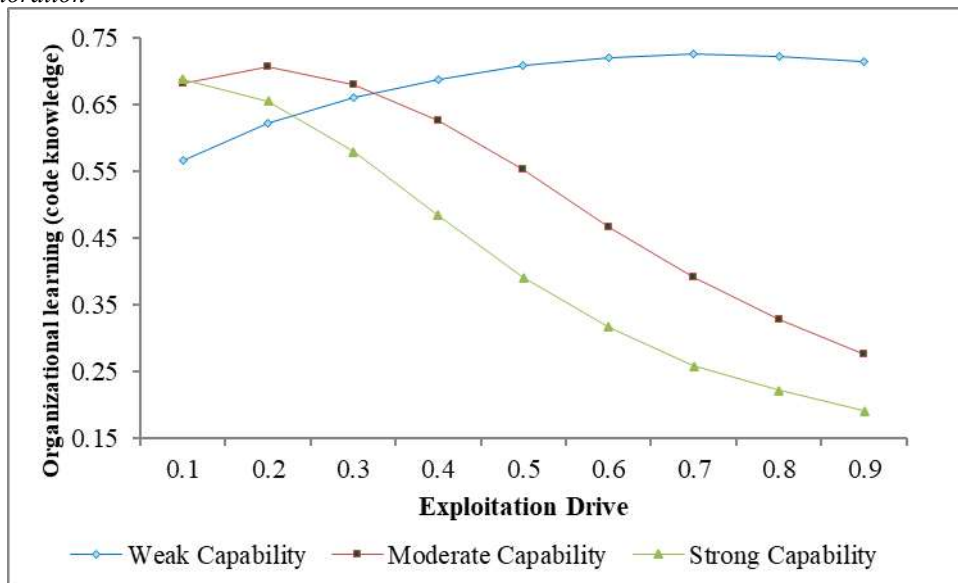



Figure 2. Variation of *organizational learning* with *exploitation drive* and *exploitation capability*, absent exploration



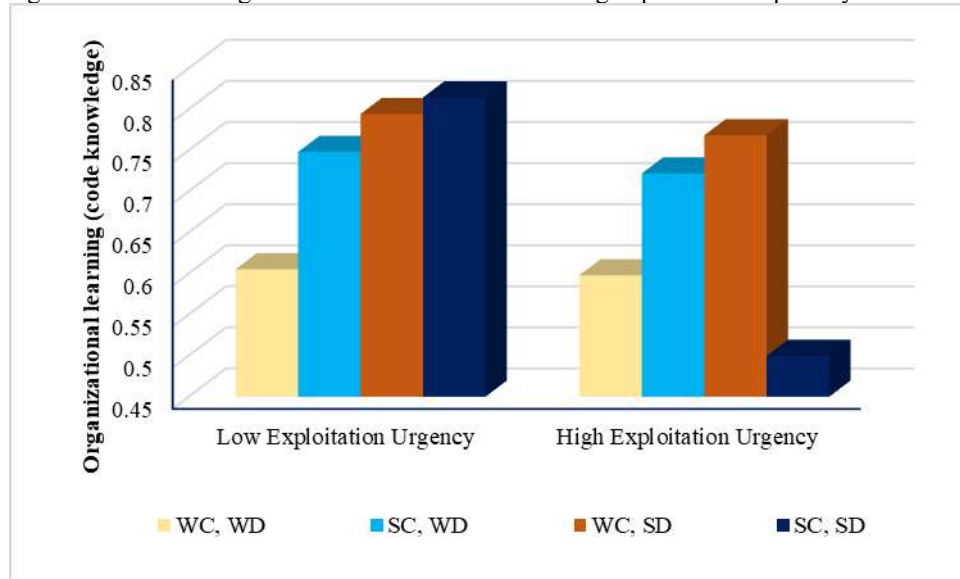
Parameters. $M = 30$; $N = 50$; Overall rate of exploitation $p_1 = 0.50$; Weak, Moderate and Strong exploitation capability correspond to 10%, 50% and 90% information processing capability respectively; Rate of exploration $p_{3_alt} = 0$; Time-steps $T = 100$.

Figure 3. Variation of *organizational learning* with *exploitation drive* and *exploitation capability*, in presence of exploration



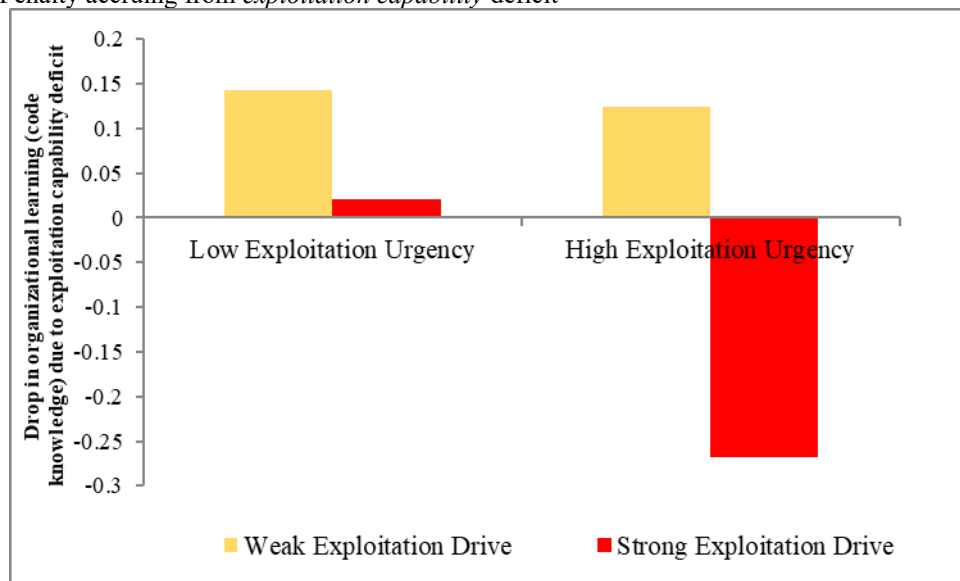
Parameters. $M = 30$; $N = 50$; Overall rate of exploitation $p_1 = 0.50$; Weak, Moderate and Strong exploitation capability correspond to 10%, 50% and 90% information processing capability respectively; Rate of exploration $p_{3_alt} = 0.10$; Time-steps $T = 100$.

Figure 4. Organizational learning outcomes under weak and strong exploitation capability and drive




Parameters. $M = 30$; $N = 50$; Overall rate of exploitation $p_1 = 0.50$; Rate of exploration $p_{3_alt} = 0.10$; Time-steps $T = 100$; Weak and Strong exploitation capability (WC, SC) correspond to 10% and 90% information processing capability respectively; Weak and Strong exploitation drives (WD, SD) correspond respectively to mobilization of 10% and 90% members for exploitation tasks in any given time-step; Low and high exploitation urgency correspond respectively to stopping of exploration after time-step 80 and 20 respectively.

Figure 5. Penalty accruing from *exploitation capability deficit*



Parameters. $M = 30$; $N = 50$; Overall rate of exploitation $p_1 = 0.50$; Rate of exploration $p_{3_alt} = 0.10$; Time-steps $T = 100$; Weak and Strong exploitation drives correspond respectively to mobilization of 10% and 90% members for exploitation tasks in any given time-step; Low and high exploitation urgency correspond respectively to stopping of exploration after time-step 80 and 20 respectively. The values on the vertical axis are obtained by subtracting the organizational learning (code knowledge) attained under weak exploitation capability from the organizational learning (code knowledge) attained under strong exploitation capability. Weak and Strong exploitation capability correspond to 10% and 90% information processing capability respectively.

Figure 6. Penalty accruing from *exploitation drive deficit*



Parameters. $M = 30$; $N = 50$; Overall rate of exploitation $p_1 = 0.50$; Rate of exploration $p_{3_alt} = 0.10$; Time-steps $T = 100$; Weak and Strong exploitation capability correspond to 10% and 90% information processing capability respectively; Low and high exploitation urgency correspond respectively to stopping of exploration after time-step 80 and 20 respectively. The values on the vertical axis are obtained by subtracting the organizational learning (code knowledge) attained under weak exploitation drive from the organizational learning (code knowledge) attained under strong exploitation drive. Weak and Strong exploitation drives correspond respectively to mobilization of 10% and 90% members for exploitation tasks in any given time-step.

APPENDIX

Section A. Case study: A European pharmaceutical company

Burgelman and Aaltonen (2024) provide an informative case study of a German pharmaceutical company (codenamed “Pharma”) where the therapeutics division attempted radical innovation in biotechnology-based molecules to cure ailments in the oncology and cardio-vascular domains over several years. Pharma did not succeed. Ultimately it was taken over by a rival. The longitudinal case study (1995-2006) surfaced certain strategic exploitation challenges typical to stage-gate projects, that inform our subsequent theory development.

In 1995, *Pharma* held a strong position in gynecology and diagnostics in Europe and Latin America. Gynecology was renowned for its wide range of hormonal contraceptive products. Diagnostics provided contrast media for X-ray and magnetic resonance imaging processes to achieve better accuracy in detecting various diseases. Besides, Pharma’s general therapeutics unit worked on cures for the central nervous system and cardiovascular diseases.

Pharma’s therapeutics division developed a new molecule for treating multiple sclerosis—a degenerative and lethal disorder of the nervous system—under the research leadership of Dr. R (a pseudonym) who was hired from a leading German University back in 1983. This treatment obtained significant market success in the US. During the mid-1990s, Pharma top management designated its therapeutic business unit as the most strategically important unit for corporate growth. Pharma had a successful public offering in the New York Stock Exchange in the early 2000s. However, this exposed Pharma to increased capital market pressure, as higher earnings and share prices were expected. Top management decided that the company had to grow faster to stay independent in the consolidating pharmaceutical industry. They undertook a two-pronged approach to accomplish this goal. Firstly, they were on the lookout for acquiring businesses with promising products. Second, Pharma top management set sights on developing competency in oncology and cardio-vascular therapeutics and signaled strategic intent by allocating up to 50% of the corporate project development budget (R&D budget) for the therapeutics division.

Pharma’s efforts in the acquisition space during the study period find somewhat sparse mention in Burgelman and Aaltonen (2024). For this reason, we focus on Pharma’s endeavor to bring about radical innovations in oncology and cardio-vascular therapeutics through in-house efforts. The following quote from Burgelman and Aaltonen (2024, p. 12) describe Pharma’s project development environment.

Pharma's stage-gate system for project development comprised five stages: a preclinical stage (A), followed by four stages (B–E) associated with regulatory-defined I–IV phases. The preclinical stage (stage A) involved conducting nonhuman studies on toxicology, carcinogenicity, pharmacology, compound stability, and determination of the manufacturing process. Phase I (stage B) involved conducting small-scale trials on healthy humans (20–50) to test pharmacology, drug effects, and dosage. Phase II (stage C) involved conducting controlled clinical trials with a moderate number of patients (30–300) to study the efficacy of the therapy. Phase III (stage D) involved conducting controlled clinical trials with various patients (300–3000) to study the safety of the therapy. Phase IV (stage E) involved conducting open clinical trials to detect possible rare side effects and to adjust for optimal drug use.

In principle, the stage-gate system was sequential, such that a project had to complete successfully previous stages before entering the next stage in the sequence. However, there were occasional exceptions in the therapeutics unit, i.e., a later phase commenced while the previous step's evaluations were still pending. In general, a project would be terminated if it failed to meet the targeted performance or if it did not fulfil a regulatory criterion in any given stage. Again, there were some exceptions for the projects in therapeutics, i.e., a terminated project would show up in a subsequent budget-cycle. In the description that follows, a move from the pre-clinical stage to later stages is considered as a move from exploration to exploitation, i.e., a move from “research” to “development”. In particular, the Stage C and Stage D—both connoting ‘exploitation’—got prominence in top management discussions on project prioritization in view of the high spending involved and the public visibility it entails. The coming into being of early-stage projects—connoting “exploration”—was subject to much lower extent of top management oversight, as described below.

Pharma's selection rule for project prioritization decisions—prospective high global (US) sales—was not applied to the early development stage where a newly created project for development of a molecule to treat an ailment enters *exploration* or research. This is in view of the difficulties associated with evaluating early promises of large potential sales in the US market. Over time, a large number of early-stage projects got prioritized in the therapeutics business unit. The Pharma top brass (that included Dr. R in every decision-making body) went along with the prioritization decisions, given Dr. R's track record of success with the multiple sclerosis cure. Such prioritization escalated the commitment of non-financial and financial resources to therapeutics. Burgelman and Aaltonen (2024, p. 15) write:

These included scientific personnel, research personnel in pharmaceutical and clinical trials involving sometimes thousands of patients worldwide, various established and new laboratory apparatus, construction of pilot production plants and their upscaling, and external subcontracted scientific and technical services.

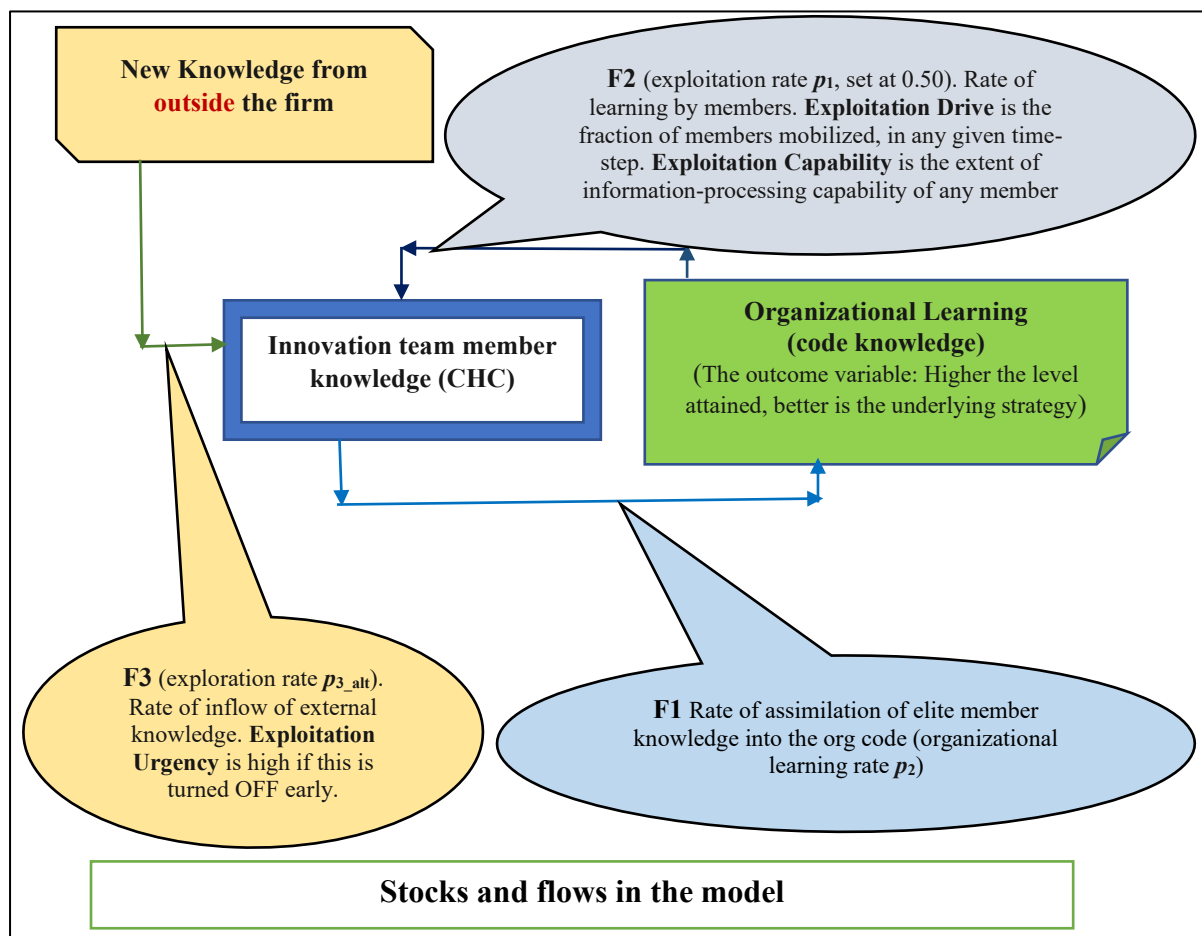
Pharma's thrust into therapeutics in the oncology and cardio-vascular domains did not succeed. Though prioritized projects in the therapeutics unit received sizeable allocations of the nonfinancial medical-related capabilities deemed necessary to speed up their progression, the number of instances of terminations of early-stage projects were significantly higher than the number of instances in the gynecology and diagnostic units. This was attributed to inexperience of managing the allocation of radically new medical-related capabilities, in the face of strong urgency set in motion by the head of corporate R&D, Dr. R. Nevertheless, certain early-stage projects made it to the development stages involving high spending and visibility, in spite of some disagreements. Around 2004, a couple of high-profile projects—one in oncology and one in cardio-vascular—failed, after expenditure of close to one billion Euro over six years. Pharma executives traced the reasons to the complicated and long-term nature of the trials, and the fact that increasing development cost and price regulation made conducting project development expensive (Grabowski and Vernon, 2000; DiMasi, 2001; DiMasi, Hansen and Grabowski, 2003).

In 2006, XYZ (a pseudonym) acquired Pharma paying 40% over Pharma's share-price. Some Pharma executives believed that the therapeutics initiatives would have succeeded if more time were provided. Others—particularly some Board members—were of the opinion that they did not see something sizeable that would give Pharma shareholders more than 40% above the stock-price. Relatedly, it is known that XYZ terminated a high-profile oncology project in the therapeutics unit shortly after acquiring Pharma.

Section B. Stocks and flows in the simulation model

In **Figure B1** we provide a schematic diagram for the model mechanics underlying the adoption of the March (1991) model pertaining to exploration through experimentation with heterogeneous external knowledge in terms of joint consideration of stocks and flows (Dierickx and Cool 1989, Bontis et al. 2002). This diagram shows the process of development of organizational learning (code knowledge)—relating back to March’s organizational learning metaphor to describe the principles of exploration and exploitation.

Figure B1. Schematic diagram for stocks and flows in the model



B.1. STOCKS AND FLOWS

In all, there are three *stock* variables and three *flow* rates connecting them.

- Stock Variable #1 is “New Knowledge”. We alternately refer to this as “heterogeneous new knowledge from outside the organization”. The source (or container) of “New Knowledge” constitutes academic and research institutions, consultants, industry trade body meets and conferences etc., consulted by the organizational members involved in exploration through experimentation.
- Stock Variable #2 is “organizational member knowledge” or “innovation team member knowledge”. It is the stock of knowledge held by organizational members involved in the

strategic innovation project (in their heads). Elsewhere (e.g., Chanda et al. 2018, von Nordenflycht 2011) this is referred to as *collective human capital* (CHC).

- Stock Variable #3 is “code knowledge” or “organizational learning” developed by the team undertaking the strategic innovation project. The innovation team owns a segment of the official organizational code to house databases, rules, norms, forms, procedures etc. associated with the strategic innovation project. Thus, “code knowledge” or “organizational learning” developed by the team resides in the “owned segment of the organizational code” (OC) distinct from the repository for broader organizational knowledge.

The *flow* variables work as follows:

- The flow **F1**—corresponding to the variable p_2 in the computational simulation model—depicts the *rate* at which the organizational code learns from mobilized members who are more knowledgeable than OC itself. Here the stock of organizational learning (code knowledge, the stock variable #3 above) gets updated from the stock of innovation team member knowledge (the stock variable #2 above)¹⁶.
- Second, the *rate* of members learning from the OC is represented by the flow **F2**, corresponding to the variable p_1 in the computational simulation model (representing exploitation). In this case the stock of organizational member knowledge gets modified by the flow **F2** from the stock representing organizational learning (code knowledge, the stock variable #3 above) developed by the team undertaking the strategic innovation project. *Exploitation Drive* is *high* if there is higher “escalation of commitment” (Burgelman and Aaltonen 2024, p. 6, p. 28) through mobilization of higher proportion of the innovation team members for this activity. *Exploitation Drive* is *low* otherwise. *Exploitation Capability*—concerning “efficaciousness in learning to master the radically new ... capabilities” (Burgelman and Aaltonen, 2024 p. 27)—is high if the concerned team member has high information-processing ability; otherwise, *Exploitation Capability* is low.
- Lastly, if exploration through experimentation is permitted, the knowledge of organizational members (stock variable #2 above) involved in the strategic innovation project gets updated (at a certain *rate* p_{3_alt}) by heterogeneous new knowledge (stock variable #1 above) from say, interactions with academic and research institutions, consultants, by members attending conferences, trade body meets and so forth. The corresponding flow is **F3**. In the simulation model, **F3** is represented by having p_{3_alt} greater than zero¹⁷. *Exploitation Urgency* is *high* if this flow is turned off early, and *Exploitation Urgency* is *low* if this flow is turned off late. In other words, “moving rapidly from exploration (research) to exploitation (development)” (Burgelman and Aaltonen 2024, p. 30) signifies high *Exploitation Urgency*, and allowing exploration (research) to continue for longer connotes low *Exploitation Urgency*.

¹⁶ This is discussed in greater detail in Section D.

¹⁷ March (1991) used a variable $p_3 = 0.10$, connoting that 100% of the innovation team members have a ten percent chance of having all their belief dimensions randomized in any given time-step. We use a lower (more realistic) rate for inflow of heterogeneous outside knowledge, where $p_{3_alt} = 0.10$, connotes that 25% of the innovation team members have a ten percent chance of having all their belief dimensions randomized in any given time-step.

B.2. MODEL MECHANICS

(I) The stock of organizational learning (code knowledge) is updated by the flow **F1** (modeled by the variable p_2) from the stock of innovation team member knowledge. Via the flow **F1**, the **OC** learns from the subset of *members carrying superior knowledge* (the ‘elites’) from among the mobilized members in the strategic innovation project team, in any given time-step. We may note that the **OC** gets correct and incorrect knowledge as input when there is correct and incorrect groupthink among the *elites*, respectively.

(II) The stock of member knowledge is updated from the stock of organizational learning (code knowledge) acquired by the team undertaking the strategic innovation project by the flow **F2** (modeled by the variable p_1 and the related constructs, *Exploitation Drive* and *Exploitation Capability*) whereby a mobilized member having a certain level of information-processing capability learns from the **OC**.

(III) When obtaining heterogeneous new knowledge from outside the organization is permitted, the stock of knowledge of members undertaking a strategic innovation project gets updated by the flow **F3** (modeled by setting the variable p_{3_alt} to a value greater than zero, for longer if *Exploitation Urgency* is low and for a shorter duration if *Exploitation Urgency* is high) from, say, interaction with industry and academic institutions, attending conferences and trade body meetings, etc. Both correct and incorrect knowledge can come in the updates to the knowledge of members of the strategic innovation project team (given that random values from $\{-1, 0, +1\}$ with one-third probability are used to update member belief dimensions).

(IV) We note that once a member learns the "generally known stuff" in the **OC**, he/she stands a greater chance of being inducted into the team of the *elites* that advise the **OC**. Thereafter, a member's unique valuable knowledge in some other dimension can help to increase the organizational learning (code knowledge) in the **OC**. Such happens by driving out incorrect knowledge in the corresponding **OC** dimension that might have been acquired earlier (due to incorrect groupthink among the the-then *elites*).

(V) In sum, the level of organizational learning (code knowledge) reached by the **OC** is a function of several parameters: (a) The extent of alignment of the initial innovation team member knowledge with the knowledge associated with the firm's current strategy. For the “new product development” / “innovation” context, we have the initial code knowledge substantially misaligned with the knowledge required for the targeted innovation¹⁸. (b) The rate of external heterogeneity infusion (p_{3_alt}) and the duration for which it is allowed to stay turned ON, which is a function of *Exploitation Urgency*. (c) The rate of exploitation (p_1) and the associated values for *Exploitation Drive* (extent of mobilization of members) and *Exploitation Capability* (information-processing capability of a given member). (d) The rate of learning by the **OC** (p_2).

¹⁸ This is implemented by creating a population deficient in Collective Human Capital, i.e., deficient in organizational member knowledge. For this first we create a *Marchian* population (where all belief dimensions have values -1, 0, or +1 with equal probability). Thereafter, we overwrite about fifteen percent of (randomly chosen) belief dimensions of members with values opposite that of the value in the corresponding dimension of the external reality (innovation target). Thus, in a dimension where the external reality has value ‘+1’, we overwrite a chosen member belief dimension with value ‘-1’, and where the external reality has value ‘-1’, we overwrite a chosen member belief dimension with value ‘+1’. Effectively, the *misalignment* is implemented by having the ratio of the initial average level of incorrect knowledge to correct knowledge in the team increase from 1:1 to 1.5:1.

B.3. EVOLUTION OF THE MODEL FROM MARCH (1991) TO BURGELMAN AND CHANDA (2024) TO THIS MANUSCRIPT

A few words are in order, regarding how the computational simulation model has evolved from March (1991) to Burgelman and Chanda (2024, forthcoming) to the research we present here. The conceptual specifications (Chanda and Miller, 2019) for the orthogonal conception of exploration and exploitation presented in March (1991, p. 79, Fig 4)—in the ‘open systems’ paradigm—is the source for Burgelman and Chanda (2024, forthcoming). In order to model a situation where a team engaging in autonomous strategic behavior attempts a *radical innovation*, Burgelman and Chanda (2024, forthcoming) lower the extent of organizational member knowledge or team knowledge—also referred to as *collective human capital* in other literature (e.g., Chanda, Ray, and McKelvey, 2018)—that the simulation experiments initiate with. This modification applies to our research as well, i.e., a far-reaching innovation is attempted in an arena where there is very low knowledge in the organization.

Moreover, in the context of autonomous strategic projects Burgelman and Chanda (2024, forthcoming) stipulate *autonomous experimentation* to comprise 25% percent of organizational members having a ten percent ($p_{3_alt} = 0.10$) probability that their knowledge-dimensions undergo randomization (signifying knowledge input from outside) in any given time-step. An implicit assumption is that personnel engaging in *autonomous experimentation* have limited time given their day jobs, and hence only a subset of them can learn from outside, in a given time-step. In the present research, the context is that of a team attempting to bring about a radical innovation in an initiative approved by top management. Here too, the characterization of exploration is that, 25% percent of organizational members are engaged in exploration in any time-step and there is a ten percent chance ($p_{3_alt} = 0.10$) that a given member’s knowledge-dimensions will be randomized. In March (1991) there is a ($p_3 = 0.10$) ten percent chance that all the belief dimensions of a member get randomized, in any given time-step—perhaps connoting an unrealistically high rate for inflow of heterogeneous knowledge from outside.

In the present research we additionally have the wherewithal to vary the timing of exploration—by invoking the parameter ‘exploitation urgency’—given that our context is that of a stage-gate process. Lastly, in March (1991) and Burgelman and Chanda (2024, forthcoming), the process of exploitation constituted *each* member learning from the organizational code about *every* dimension of knowledge, in *each* time-step, with probability p_1 . In our work, only a fraction of members actively take part in p_1 and p_2 processes, in any

given time-step. This is modeled as ‘exploitation drive’, indicating the extent of mobilization of personnel. Moreover, ‘exploitation capability’ is modeled as the product of the fraction of knowledge dimensions for which p_1 learning is carried out, for a fraction of the time-steps. This parameter, therefore, signifies the extent of information processing capability. The parameter p_1 continues to be the probability that a knowledge-dimension is learned by a mobilized member, when the information-processing capability allows acting on that knowledge-dimension for learning in a particular time-step. Table **AT1** shows a summary.

Table AT1. Model evolution at a glance

<i>Parameter</i>	March (1991) [M91]	Burgelman & Chanda (2024) [BC2024]	This manuscript
Rate of Exploitation (p_1)	Probability that <i>any</i> bit is learned (from the org code) by <i>any</i> member, in <i>any</i> time-step.	Same as M91	Probability that a mobilized member learns about a particular bit in a given time-step. Extent of exploitation capability (information-processing ability) determines whether a given bit is up for update in a given time-step.
Rate of learning by code (p_2)	Related to the rate at which the code learns from elite members	Same as M91	Same as M91
Rate of exploration (p_3)	Probability that all bits of a member get randomized, in any time step	Probability that all bits of any randomly-chosen 25% fraction of members get randomized, in any time step.	Same as BC2024
Initial team knowledge (CHC)	Member strings are constructed by randomly filling “1”, “0”, and “-1” with equal (one-third) probability	After generating the initial population as in M91, 15% of member bits are overwritten with values opposite to that in the corresponding dimension of the reality string.	Same as BC2024
Duration of exploration	Throughout all time-steps	Same as M91	Limited to few time-steps according to the level of exploitation urgency

Section C. Model-based explanation of selected results

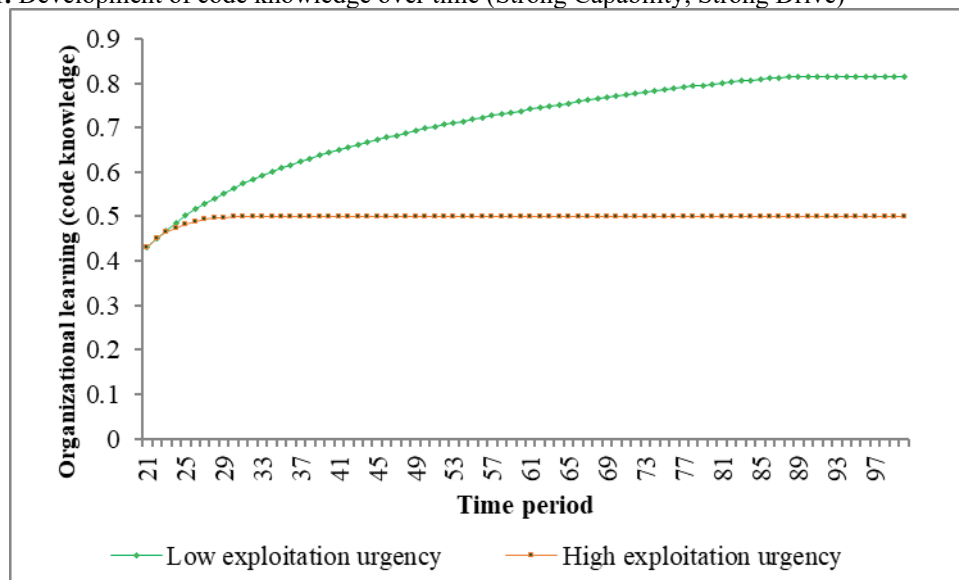
This section presents explanations for the selected results as follows:

- First, we explain why the outcomes for (Strong Capability, Strong Drive) [SCSD] are higher under low exploitation urgency, compared to that under high exploitation urgency.
- Next, we take up explaining the rank order of outcomes under high exploitation urgency.
- Thereafter, we explain the rank order of outcomes under low exploitation urgency.

C.1. COMPARISON OF SCSD OUTCOMES UNDER LOW AND HIGH EXPLOITATION URGENCY

In **Figure C1** we show how the outcome variable *organizational learning* [code knowledge] shapes up over time, for the (Strong Capability, Strong Drive SCSD) cases, under low and high exploitation urgency. Given that the curves develop the same way in the first 20 time-steps when exploration is turned ON in both cases, we focus on the period from time-step 21 to time-step 100. We recall that, in this latter period, exploration is turned OFF throughout, for the high urgency case; for the low urgency case, exploration gets turned off after time-step 80.

Figure C1. Development of code knowledge over time (Strong Capability, Strong Drive)



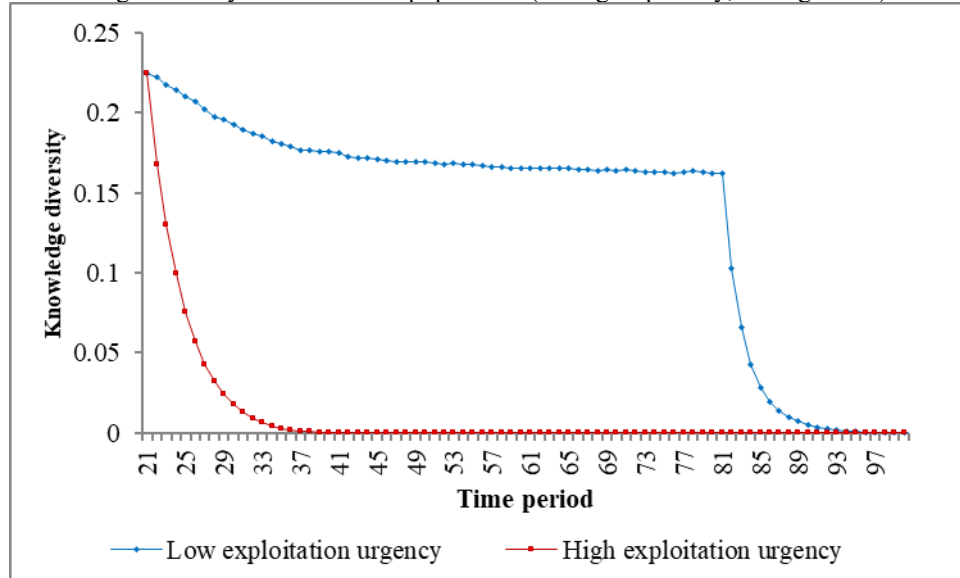
Parameters. $M = 30$; $N = 50$; Overall rate of exploitation $p_1 = 0.50$; Rate of exploration $p_{3_alt} = 0.10$ when turned ON, 0 otherwise; Time-steps $T = 100$; Strong exploitation capability SC corresponds to 90% information processing capability; Strong exploitation drive SD corresponds to mobilization of 90% members in any given time-step; Low and high exploitation urgency correspond respectively to stopping of exploration at time-step 80 and 20 respectively.

We observe that, for the curve for ‘High exploitation urgency’ the code knowledge stops developing at around the time-step 40. This happens because high mobilization in conjunction with high rate of information processing quickly drives out the variety of knowledge among members. When all members reach the level of knowledge matching the knowledge of the organizational code, the latter has no one to learn from. Hence growth in code knowledge stops. In contrast, the curve ‘Low exploitation urgency’ shows continued development of code knowledge, and flattening is noticed shortly after exploration is turned off after time-step 80. The supply of diverse knowledge arising from longer exploration is responsible for continued growth in code knowledge.

In order to confirm above intuition, in **Figure C2** we present graphs for knowledge diversity in the member population over time, for the situation depicted Figure C1. We use Shannon’s formula

for information entropy to compute knowledge diversity. For a given knowledge dimension, first we compute the proportions of “-1”, “0”, and “+1” across all the members. Considering these proportions as probability of occurrence, we compute information entropy for a given knowledge dimension as the summation of the proportion multiplied with the logarithm (base 3) of the reciprocal of the proportion. Thereafter, knowledge diversity is computed by averaging information entropy over all the knowledge dimensions.

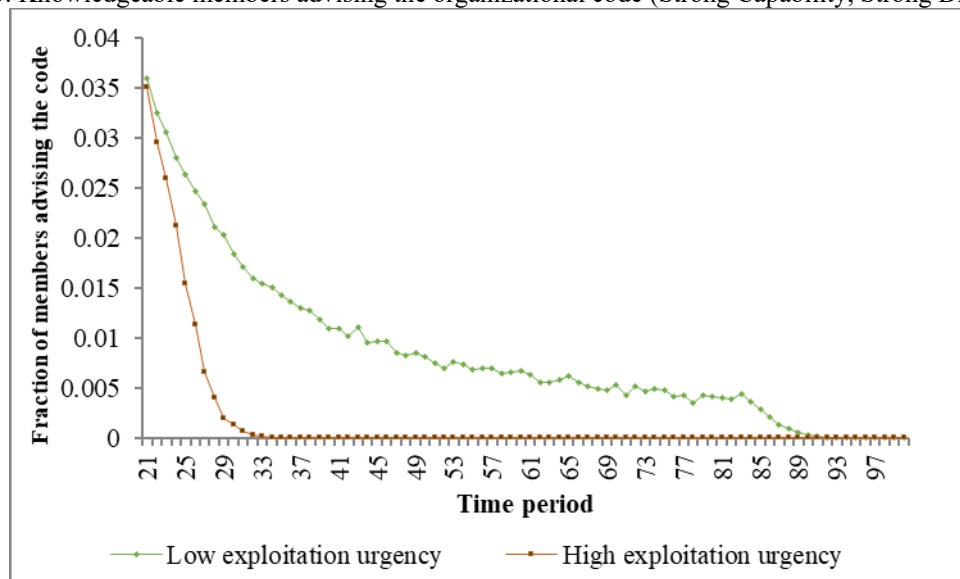
Figure C2. Knowledge diversity in the member population (Strong Capability, Strong Drive)



Parameters. Same as in Figure C1.

We learn from Figure C2 that for ‘High exploitation urgency’, diversity of knowledge in the member population falls to zero around time-step 40. This explains why code knowledge stops increasing, as seen in Figure C1. Moreover, Figure C2 also shows that for ‘Low exploitation urgency’, there is ample knowledge diversity in the member population for far longer. This enables code knowledge to keep increasing. Thereby for a conjunction of strong exploitation capability and strong exploitation drive obtains superior organizational learning outcomes under low exploitation urgency.

Figure C3. Knowledgeable members advising the organizational code (Strong Capability, Strong Drive)



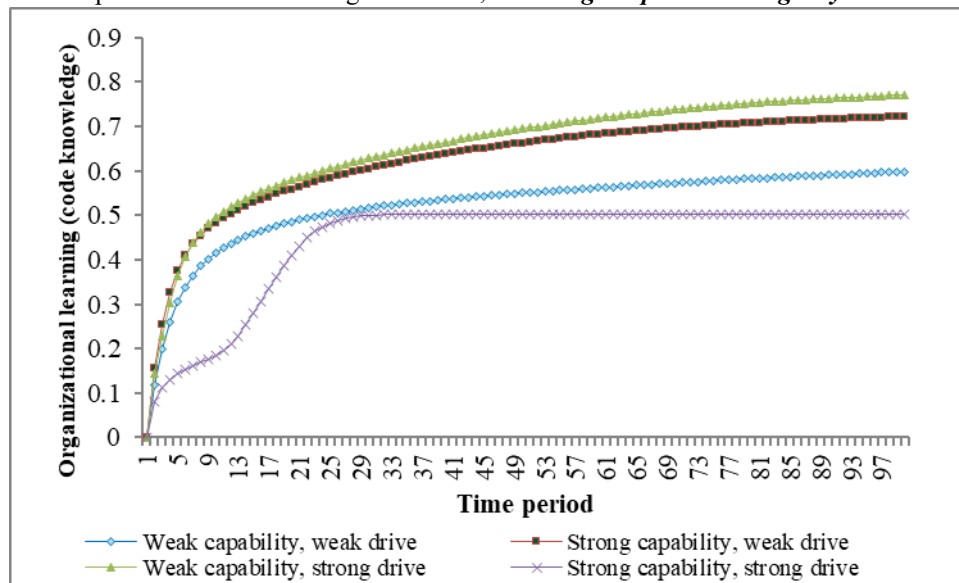
Parameters. Same as in Figure C1.

In **Figure C3**, we present the fraction of members who are more knowledgeable than the organizational code, comprising the set of members who advise the code in a given time-step. We capture this value programmatically in each time-step at the point that members who are more knowledgeable than the code are identified, in order to be deployed for the p_2 process. Again, we observe that this fraction drops to zero for the curve ‘High exploitation urgency’ at around time-step 40, confirming that the code has no one to learn from, when knowledge diversity in the member population drops to zero (as seen from **Figure C2**). In contrast, the curve ‘Low exploitation urgency’ in **Figure C3** continues to have knowledgeable members to advise the code till an advanced stage, owing to continued presence of knowledge diversity in the member population (**Figure C2**).

C.2. COMPARISON OF OUTCOMES UNDER HIGH EXPLOITATION URGENCY

In **Figure C4**, as before, we observe that, for the curve ‘Strong capability, strong drive’ (**SCSD**), the code knowledge stops developing at around the time-step 40. In contrast, for the other curves, the code knowledge continues to develop till the end of the observation period. This shows that all these other curves—including the curve ‘Weak capability, weak drive’ (**WCWD**)—obtain higher outcomes than **SCSD**, under high exploitation urgency.

Figure C4. Development of code knowledge over time, under *high exploitation urgency*



Parameters. $M = 30$; $N = 50$; Overall rate of exploitation $p_1 = 0.50$; Rate of exploration $p_{3_alt} = 0.10$ when turned ON, 0 otherwise; Time-steps $T = 100$; Weak and Strong exploitation capability (WC, SC) correspond to 10% and 90% information processing capability respectively; Weak and Strong exploitation drives (WD, SD) correspond respectively to mobilization of 10% and 90% members in any given time-step. High exploitation urgency corresponds to stopping of exploration at time-step 20.


As noted earlier, for the curve **SCSD**, high mobilization (exploitation drive) in conjunction with high rate of information processing (exploitation capability) quickly drives out the variety of knowledge among members, leaving the org code with no one to learn from, around time-step 40. In contrast, for the other curves, learning by the organization persists for longer, since the variety in member knowledge is preserved for longer, given lower rates of mobilization and/or information processing. For example, for the **WCWD**, member knowledge variety is preserved for longer, on account of the meagre rates of information-processing (**WC**) and mobilization (**WD**). In this regard it

is instructive to recall that, in this latter period, exploration is turned OFF throughout, making the knowledge-variety in the heads of members the only possible source for enhancing code knowledge.

Figure C5 shows the percentage of elite members—members who have more knowledge than the code—advising the organizational code after time-step 40, at which point the **SCSD** ceases to have further change to code knowledge given that there are no elite members. The graphs presented in **Figure C5** confirm that for curves other than the **SCSD**, a non-zero number of elites can be found in the period after time-step 40, given lower mobilization or lower rate of information processing. In this case, the increase in code knowledge (organizational learning) takes place by utilizing the diverse knowledge residing in the heads of members.

Moreover, in **Figure C5** the rank order is ‘Weak capability, strong drive’ (**WCSD**), followed by ‘Strong capability, weak drive’ (**SCWD**), and then by **WCWD** and finally the **SCSD**. This rank order is reflected in the results presented in **Figure 4** (right panel) of the main manuscript. We observe that **SCWD** has lower proportion of members that are elite, compared to **WCSD**. This is the case because lesser numbers of personnel are mobilized, in the former case.

Figure C5. Extent of elite members advising the organizational code, under *high exploitation urgency*



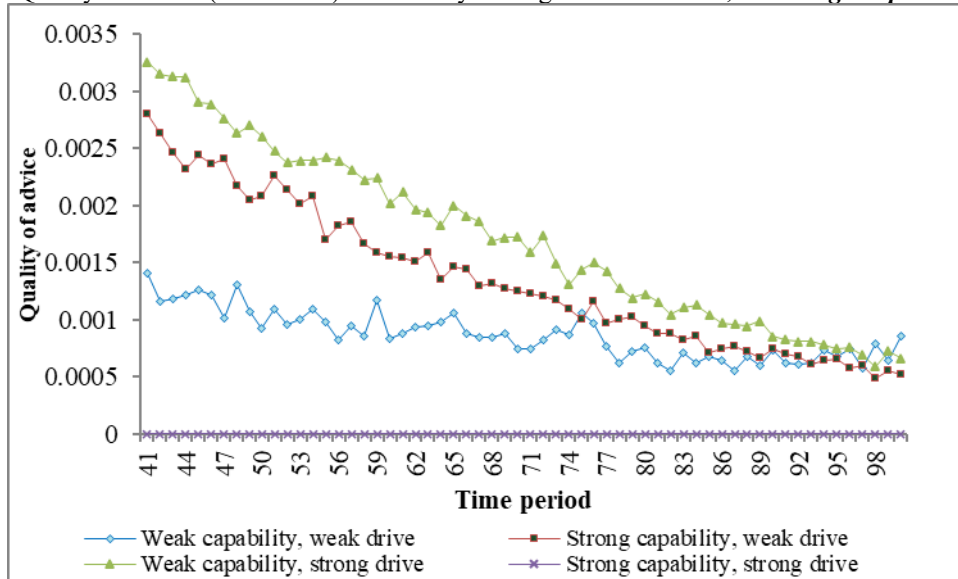
Parameters. Same as in **Figure C4**.

Next, we check whether merely having more elites advising the code leads to superior code knowledge or an additional criterion is necessary, that the recommendations from the elites have to be of good quality as well. In **Figure C6** we show this information.

We define quality of advice as the net correct suggestions [the number of correct suggestions minus the number of incorrect suggestions] / [number of knowledge dimensions] from the knowledgeable members that got implemented (i.e., incorporated into the org code). We don’t give credit or discredit if the elites recommend a value that is already present in the organizational code. Further, if the p_2 process fails to incorporate a recommendation from the elite, no credit or discredit is assessed.

We observe that **Figure C6** also shows the same rank order as found in the left panel of **Figure 4** in the main manuscript. From **Fig C6** and **C5** we infer that the rank order of outcomes is indeed a joint outcome of higher number of elites and higher extent of superior-quality recommendations.

Figure C6. Quality of advice (from elites) received by the organizational code, under *high exploitation urgency*



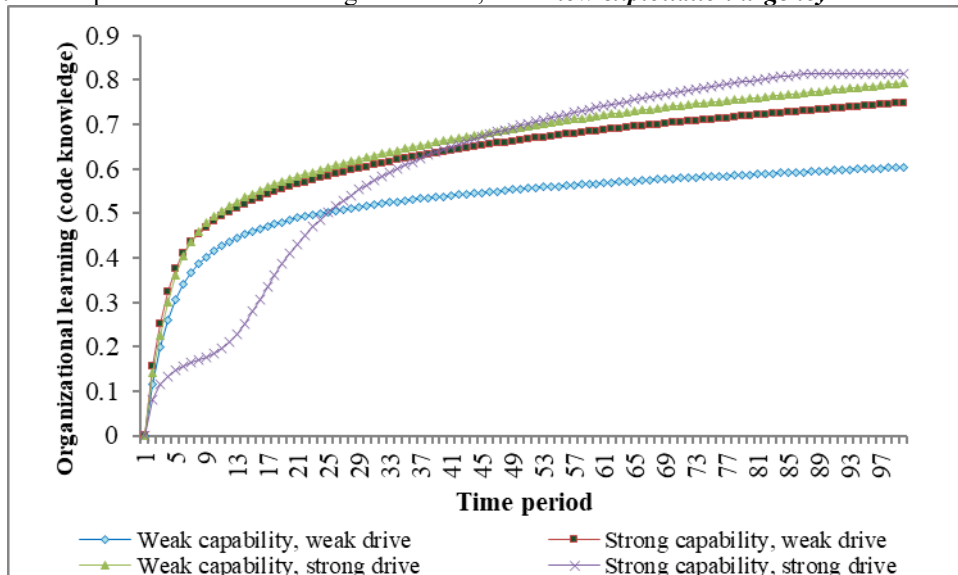
Parameters. Same as in Figure C4.

In sum, the **SCSD** curve attains the worst outcomes under high exploitation urgency because it quickly runs out of knowledge diversity. The **WCSD** attains the best outcomes because it has more elites giving comparable quality of advice as **SCWD**. The **WCWD** comes third, having less elites who provide slightly inferior advice.

C.3. COMPARISON OF OUTCOMES UNDER LOW EXPLOITATION URGENCY

In **Figure C7**, we observe the development of code knowledge over time underlying the bar-graphs presented on the left panel of Figure 4 of the main manuscript, i.e. for low exploitation urgency.


Figure C7. Development of code knowledge over time, under *low exploitation urgency*



Parameters. $M = 30$; $N = 50$; Overall rate of exploitation $p_1 = 0.50$; Rate of exploration $p_{3_alt} = 0.10$ when turned ON, 0 otherwise; Time-steps $T = 100$; Weak and Strong exploitation capability (WC, SC) correspond to 10% and 90% information processing capability respectively; Weak and Strong exploitation drives (WD, SD) correspond respectively to mobilization of 10% and 90% members in any given time-step; Low exploitation urgency corresponds to stopping of exploration at time-step 80.


In **Figure C8A** we present the fraction of members advising the organizational code, under low exploitation urgency; **Figure C8B** shows an amplified version of the situation after time-step 40. We find that the curve **SCSD** is having by far the most elites providing advice to the code for a long time, and the curve **WCWD** has the least. The strong capability and drive allow utilizing more of the heterogeneous knowledge coming from exploration. The two middle positions are taken by **WCSD** and **SCWD**, respectively, reflecting the rank order in the left panel of Figure 4 in the main manuscript.

Figure C8A. Extent of elite members advising the organizational code, under *low exploitation urgency*



Parameters. Same as in Figure C7.

Figure C8B. Extent of elite members advising the organizational code, under *low exploitation urgency*

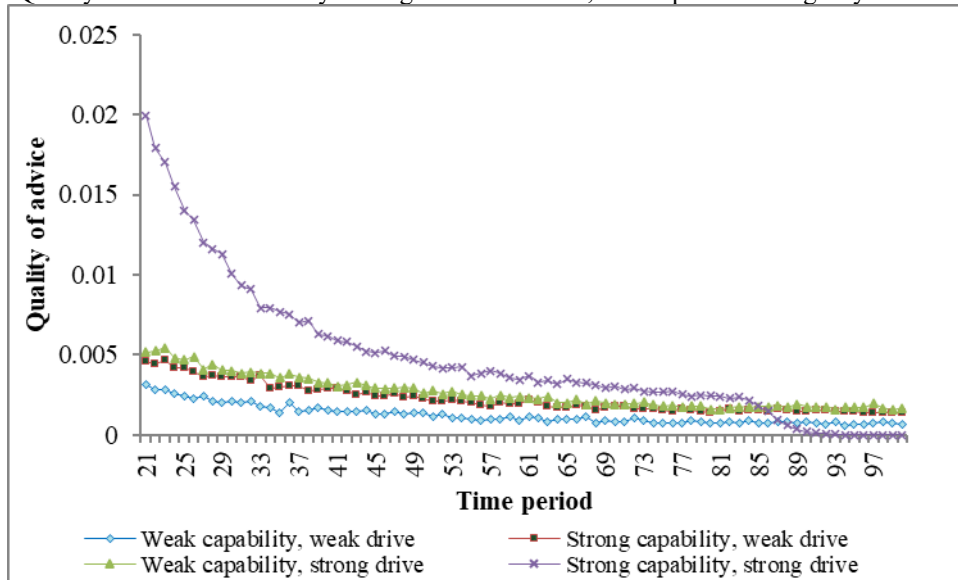


Parameters. Same as in Figure C7.

In **Figure C9**, we present curves representing the quality advice received by the organizational code, under low exploitation urgency. Here we focus on the period after time-step 20, where fortunes are different under low and high exploitation urgency. We observe that the curve **SCSD** gets higher quality of advice than the others. Thus, the combination of strong drive (high mobilization) and strong capability (high extent of information-processing) leads to higher extent of elites providing superior

quality of advice. Hence **SCSD** attains the most superior level of code knowledge. The rank order observed in the left panel of Figure 4 [**SCSD** – **WCSD** – **SCWD** – **WCWD**] is seen to reflected here as well.

Figure C9. Quality of advice received by the organizational code, low exploitation urgency



Parameters. Same as in Figure B7.

In Section **D** of the Appendix that follows, we present additional results for very low values of the variable p_2 (rate of code learning). We are presenting these results since there are minor deviations of the propositional finding from our main results. However, since an organization with low value for p_2 is unlikely to survive, these deviations are unlikely to be of practical salience.

Section D. Additional inquiry with extreme values of the organizational learning rate (p_2), a parameter we keep fixed in all our experiments

D.1. INTRODUCTION

The rate of organizational learning (p_2) is kept constant at 0.50 in all our reported results, following the lead from March (1991), Figures 2 ... 5 (pp. 77–80). In experiments to carry out robustness checks, we confirmed that the propositions we developed hold good for p_2 set at 0.40 and 0.60. However, we were curious to know if they hold good for wider range of values of p_2 . We reasoned that, if p_2 were to be zero, the organization will develop zero code knowledge (organizational learning); in such case, the propositions will fail to hold good. In additional experiments to check this intuition, we found that the propositions hold good for larger values of p_2 (e.g., for $p_2 = 0.90$); for a small value of p_2 , (0.10) we noticed minor deviations from the propositions reported. It is likely that an organization with low value for p_2 is unlikely to survive; hence these deviations are unlikely to be of practical salience. We are including the results here for the sake of completeness.

D.2. RELATING TO THE CONSTRUCT ‘ABSORPTIVE CAPACITY’

Before we present the additional results, a discussion relating the construct for the rate of organizational learning (p_2) to the construct ‘*absorptive capacity*’ at an *organizational level* proposed by Cohen and Levinthal (1990), (hereafter **CL1990**) is warranted. In this discussion, all the features and processes of ‘our’ model that we allude to applies to any genetic algorithm model drawing from March (1991) and involving the p_2 construct. On ground of comprehensiveness, we further include a discussion relating other model constructs to the construct of *absorptive capacity at the level of an individual* appearing in CL1990.

CL1990 (p. 128) define ‘**absorptive capacity**’ as the ability of a firm to “recognize the value of new, external information, assimilate it, and apply it to commercial ends” and suggest that “it is largely a function of the firm's level of prior related knowledge”. CL1990 introduce two terms: (a) ‘absorptive capacity’ at an individual level, (b) ‘absorptive capacity’ at an organizational level.

D.2.1. COMPARISON FOR INDIVIDUAL LEVEL ABSORPTIVE CAPACITY

In the formal model presented in this manuscript, ‘*absorptive capacity*’ at an *individual level* could be mapped to the construct for rate of member learning (p_1) with certain qualifications as noted below.

- The code knowledge (organizational learning) is **put to use for commercial ends**, not the member knowledge directly. It just so happens that if equilibrium conditions are reached, the member knowledge and the code knowledge are coincident.
- Our present model does **not** have a direct mechanism that makes p_1 a **function of the extent of prior knowledge possessed** by an individual member.
- The construct p_1 involves learning from the knowledge in the organizational code (**OC**), i.e., obtaining **internal (rather than external)** knowledge.
- A member can obtain **both** correct (**valuable**) and incorrect (**not valuable**), knowledge from the code, since the code can have both kinds of knowledge.
- If exploration is turned ON, a member of an innovation team has a second mechanism for obtaining knowledge—from outside the organization (say by attending conferences, industry trade body meetings, collaborating with educational / other institutions for research, etc.). The rate of exploration (p_{3_alt}) could be related to a **second kind of** absorptive capacity at an individual level, an external-facing kind— as different from the internal-facing kind when learning from the organizational code.

D.2.2. COMPARISON FOR ORGANIZATIONAL LEVEL ABSORPTIVE CAPACITY

The construct for the rate of organizational learning (p_2) could be mapped to the construct (b), in CL1990, i.e., 'absorptive capacity' at an organizational level, with certain qualifications.

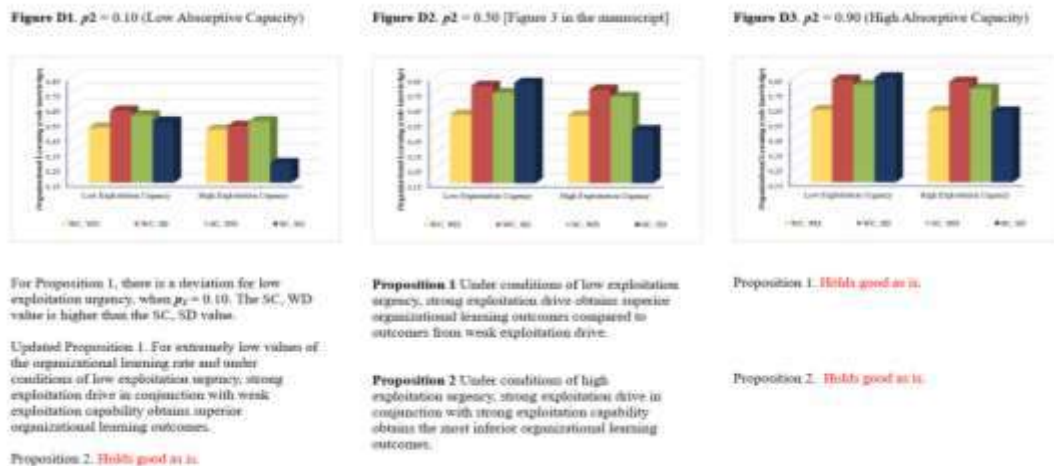
- The construct p_2 is definitely related to **assimilation** of information—availed from members of the innovation team—into code knowledge (organizational learning).
- The code knowledge developed by an innovation team *can* be conceived as the relevant construct that is **put to use for commercial ends**—given that we determine the relative worth of alternate strategies by choosing organizational learning (code knowledge) as the outcome variable, representing *fitness*.
- Moreover, p_2 is indeed a function of **prior code knowledge**, since the elites that are chosen to advise the organizational code are those members who have more knowledge than the code¹⁹.
- A process of **recognizing the value of new information** also exists in the p_2 process, since the code learns from the elites through a majority-vote mechanism. We note though, in our model, even incorrect information may be erroneously recognized as valuable—meriting emplacement as code knowledge (organizational learning)—if the majority of the elites have an incorrect belief in the concerned knowledge dimension. CL1990 (p. 138) recognize this point somewhat differently, that a firm may fail to accord 'value' to a 'valuable' piece of information on account of lacking the capability or absorptive capacity to do so. The following quote is relevant here: “firms may not realize that they should be developing their absorptive capacity due to an irony associated with its valuation: the firm needs to have some absorptive capacity already to value it appropriately.” (p. 138). The tone of CL1990 expresses a yearning for a kind of *dynamic capability*—i.e. higher order capability (Winter 2003)—that enables recognizing a piece of information as valuable or otherwise, and an admission that such may not be feasible (to develop) in environments rife with uncertainty.
- Further, CL1990 (p. 128) relate 'absorptive capacity' at an organizational level to recognizing “the value of new, **external** information” (boldface and underlining of the word 'external' ours). In our model, the construct for the rate of organizational learning (p_2) is involved in recognizing the value of *any* information, external (obtained by inflow of heterogeneous knowledge from outside the organization when exploration is turned ON) or internal (knowledge originally possessed by members), that happens to reside in the heads of team members.

D.3. FORMAL MODELING RESULTS WITH EXTREME VARIATION OF p_2

In this section, we furnish results of experiments underlying Figure 4, 5 and 6 of the main manuscript by running those experiments with p_2 set at 0.10 (leftmost panel) and at 0.90 (rightmost panel). For ease of comparison, we show the manuscript results with p_2 set at 0.50 in the middle panel.

¹⁹ The elite selection mechanism is equivalent to our cultivator-ancestors keeping aside the stronger seeds for sowing for the next year, without knowing precisely what made those seeds 'stronger' / 'fitter' to reproduce. The fact that humanity—still dependent on cultivation and plantation as primary or secondary food source—has survived, is testimony to the wisdom of the decisions by our cultivator-ancestors.

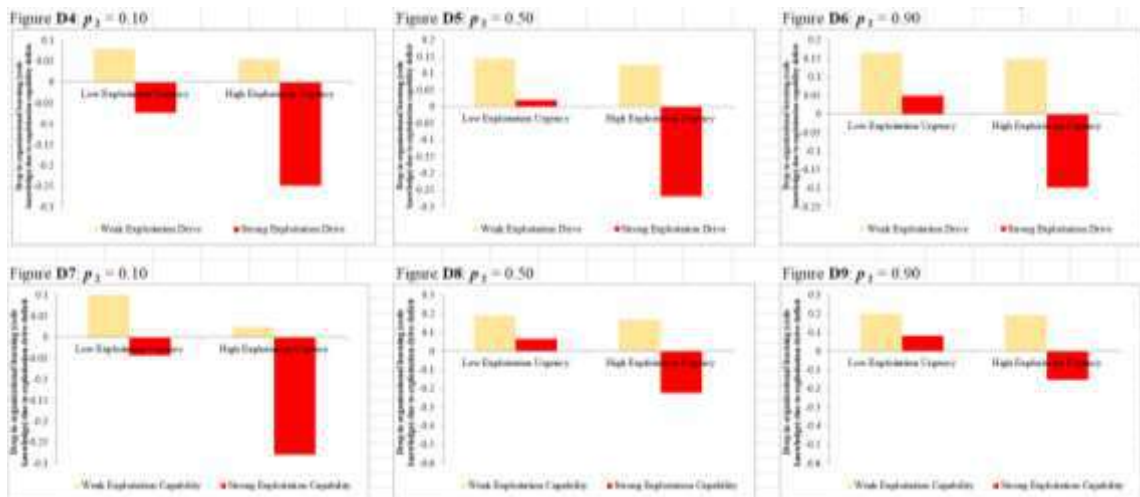
D.3.1. COMPARISON FOR PROPOSITIONS 1 & 2 OF THE MANUSCRIPT



Comparing the results shown in Figures D1, D2 and D3 we infer that:

- (a) Proposition 2 (under conditions of high exploitation urgency, strong exploitation drive in conjunction with strong exploitation capability obtains the most inferior organizational learning outcomes) holds good *as is* in all the cases and
- (b) for low values of p_2 , Proposition 1 (under conditions of low exploitation urgency, strong exploitation drive obtains superior organizational learning outcomes compared to outcomes from weak exploitation drive) could be framed with a qualification as follows: *For extremely low values of the organizational learning rate and under conditions of low exploitation urgency, strong exploitation drive in conjunction with weak exploitation capability obtains superior organizational learning outcomes compared to outcomes from weak exploitation drive.*

D.3.2. COMPARISON FOR PROPOSITION 3 OF THE MANUSCRIPT



Comparing the results shown in Figures D4, D5 ... D9 we infer that Proposition 3 (under conditions of high exploitation urgency, organizational learning outcomes are not penalized by exploitation drive deficit if exploitation capability is strong; or by exploitation capability deficit if exploitation drive is strong) holds good *as is* in all the cases.

Section E. Similarities and differences between the genetic model from March (1991) that we invoke, and Kauffman's (1993) *NK* model and the Bandit model that we do not invoke (but could be more familiar to scholars of management and organization studies)

Kauffman's *NK* model (e.g., Knudsen and Levinthal, 2007)—as well as an *N*-armed bandit model (e.g., Posen and Levinthal, 2012)—usually appear as a **single agent** model²⁰ involved with the task of finding values in *N* cells of an array that confer the maximum attainable *fitness* or *payoff*, given constrained resources. In contrast the March (1991) genetic model allows observation of macro-level outcomes from stylized micro-level actions of **multiple agents**. In the single-agent framing, bandit and *NK* models do not have a provision for considering individual *and* organization level absorptive capacity (please refer to our discussion in Section D.2 of the Appendix) separately; this shortcoming is not present in March's genetic model. The latter also allows distinct characterization of knowledge heterogeneity of organizational members (such is not possible in single-agent models). Thus, March's model is suitable for consideration of detailed constructs like *exploitation drive*, *exploitation capability* and *exploitation urgency*, in the context of a strategic innovation project.

SIMILARITIES: *NK* MODEL (KAUFFMAN), GENETIC MODEL (MARCH) AND BANDIT MODEL

[I]. The models involve algorithmically populating an *N*-bit (*N*-cell) string (array), the '*decision-configuration*' or '*outcome*' variable. Each cell or bit may be considered as a *decision-variable*, and the values obtained in the *N* cells—upon running specific algorithms— determine '*fitness*' / '*payoff*'.

[II]. *NK* and Genetic models aim to attain '*decision-configurations*' having as high '*fitness*' as feasible. In a Bandit model, an important objective is to identify pay-off probability for a given arm.

[III]. The models involve running respective computational simulation experiment (algorithm) over the duration of certain number of time-periods, where the '*outcome*' gets refined over time.

[IV]. Allowable cell values are "+1" and "-1" in case of March, and "0" and "1" in case of Kauffman; in either case, we could use "A" and "B" instead, without any loss of generality²¹. For a Bandit model, a "1" signifies that an arm was activated in quest of a payoff, and "0" signifies non-activation.

[V]. The models yield insight(s) regarding '*adaptation*' by an organism or by an organization.

- In case of Kauffman (1993), a key insight is that adaptation becomes very difficult at high complexity²², given by the extent of interdependence among the '*decision-variables*'.
- In case of March (1991), a key insight is that faster adaptation can result in lower overall fitness (Figure 1, p. 76 and Figure 2, p.77). A second insight is that in an open system, the fitness-destroying effect of environmental turbulence (or disequilibrium conditions) can be successfully countered by (presence of) a mutation-like mechanism for inflow of heterogenous knowledge from outside the system (Figure 5, p. 80).
- A Bandit model is used for checking efficacy of various mechanisms to find payoff probabilities

²⁰ Some exceptions are: two-agent models in Rivkin and Siggelkow (2003), Siggelkow and Rivkin (2005) etc.

²¹ March's model allows a third value "0" as a cell value for a *decision-variable*, signifying '*no opinion*'. However, the value "0" rarely appears in the output/outcome array from a computational simulation experiment. We could use the alphabet "C" in place of "0" without any loss of generality.

²² Kauffman labels this phenomenon "*complexity catastrophe*". See also Chanda and Yayavaram (2021).

In Table AT2 that follows we present the differences between the models.

Table AT2. Contrast between the NK, Genetic and Bandit models

Feature	<i>NK model</i>	<i>March's genetic model</i>	<i>Bandit model</i>
<i>Origination / Vintage</i>	Kauffman's adaptation of randomized binary networks in the late 1980s, early 1990s	March's (1991) adaptation of Holland's work (1975) on genetic algorithms	Finds mention in Simon's work in the 1950s
<i>Number of agents</i>	Usually, single agent	Multi-agent	Single Agent
<i>Study of emergence (micro to macro transformation)</i>	Not feasible for single-agent model	Feasible. Allows study emergence of macro-level pattern from stylized micro-level mechanisms.	Not feasible
<i>Key mechanism(s)</i>	Fitness-walk according to a stylized logic	Variation-selection-retention, Crossover and Mutation (genetic mechanisms)	Bayesian (or other) updating of probabilities
<i>Interdependence among decision dimensions</i>	Varying interdependence maps to varying complexity. Value in a cell, in conjunction with values in other cells having dependency with the focal cell determines fitness contribution	Not applicable.	Usually not applicable
<i>Learning</i>	None. The decision-maker can determine whether a change is acceptable by computing the change in fitness.	Learning over time is involved. Some incorrect knowledge gets developed as well.	Learning through Bayesian updating (or equivalent)
<i>Key reason of variability of outcomes</i>	Parameter settings that enable more thorough search (e.g., inspecting a higher number of configurations), or entailing lower interdependence ²³ yield high fitness	Parameter settings that enable longer persistence of diversity of knowledge of organizational members usually yield higher fitness ²⁴	Type of imperfection in actions based on imperfect knowledge of underlying probabilities
<i>Reason for using computational simulation (Mathematical tractability)</i>	Mathematical tractability is not a concern. Computational simulation model is used only because, for large "N", large time is required to compute fitness of 2^N configurations. We try to obtain the best we can, given limited number of computations.	Mathematical tractability is low. Computational simulation is used because crossover and mutation are hard to model by mathematics. The model allows for stochastic and unpredictable change at run-time.	Mathematical tractability is not a concern. Payoffs tend towards underlying probabilities as they get determined with more accuracy.
<i>Unpredictable change / Open systems use cases</i>	Not suited. The fitness landscape is generated, one-time, at the beginning and remains the same throughout ²⁵ .	Can be studied by enabling the mechanisms for (a) turbulence, that changes the external reality over time and (b) mutation that update agents' characteristics	Not suited. Under turbulence, outcomes tend to converge towards random choice.

²³ Exception being interdependence level increasing from 1 to 3, where a peak is encountered.

²⁴ An exception being Figure 4 in March (1991, p. 79) the curve $p_1 = 0.1$.

²⁵ Changing the fitness landscape several times during the course of the simulation tends to lower the average fitness outcome to around 0.50; the same can be obtained by using a coin-toss.

References to the Appendix

- Burgelman RA, Aaltonen P (2024, *forthcoming*) Fading corporate survival prospects: Impact of co-selection bias in resource allocation on strategic intent. *Strategic Management Journal*.
- Burgelman RA, Chanda SS (2024, *forthcoming*) Autonomous strategic behavior, organizational learning and top management support: Re-examining field research with computational modeling. *Strategic Management Review*.
- Bontis N, Crossan MM, Hulland J (2002) Managing an organizational learning system by aligning stocks and flows. *Journal of Management Studies* 39(4):437–469.
- Chanda SS, Ray S, McKelvey B (2018) The continuum conception of exploration and exploitation: An update to March's theory. *Management Science* 64(3):1050–1079.
- Chanda SS, Yayavaram S (2021) *Overcoming Complexity Catastrophe: An Algorithm for Beneficial Far-Reaching Adaptation under High Complexity*. arXiv: <http://arxiv.org/abs/2105.04311>.
- Cohen WM, Levinthal DA (1990) Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly* 35(1):128–152.
- Dierickx I, Cool K (1989) Asset accumulation and sustainability of competitive advantage. *Management Science* 35(12):1504–1511.
- DiMasi JA (2001) Risks in new drug development: Approval success rates for investigational drugs. *Clinical Pharmacology and Therapeutics* 69, 297–307.
- DiMasi JA, Hansen RW, Grabowski HG (2003) The price of innovation: new estimates of drug development costs. *Journal of Health Economics* 22, 151–185.
- Grabowski H, Vernon J (2000) The determinants of pharmaceutical research and development expenditures. *Journal of Evolutionary Economics* 10, 201–215.
- Holland JH (1975) *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: University of Michigan Press.
- Kauffman SA (1993) *The Origins of Order: Self-organization and Selection in Evolution*. Oxford, UK: Oxford University Press.
- Knudsen T, Levinthal DA (2007) Two faces of search: alternative generation and alternative evaluation. *Organization Science* 18(1):39–54.
- March JG (1991) Exploration and exploitation in organizational learning. *Organization Science* 2(1):71–87.
- Posen HE, Levinthal DA (2012) Chasing a moving target: exploitation and exploration in dynamic environments. *Management Science* 58(3):587–601.
- Rivkin JW, Siggelkow N (2003) Balancing search and stability: interdependencies among elements of organizational design. *Management Science* 49(3):290–311.
- Siggelkow N, Rivkin JW (2005) Speed and search: designing organizations for turbulence and complexity. *Organization Science* 16(2):101–122.
- 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.
- Winter S (2003) Understanding dynamic capabilities. *Strategic Management Journal* 24(10):991–995.