

BALANCING EXPLORATION AND EXPLOITATION  
IN BOTTOM-UP ORGANIZATIONAL LEARNING CONTEXTS

By

Ross Ian Walker

A THESIS

Submitted to  
Michigan State University  
in partial fulfillment of the requirements  
for the degree of

Psychology – Master of Arts

2018

## ABSTRACT

### BALANCING EXPLORATION AND EXPLOITATION IN BOTTOM-UP ORGANIZATIONAL LEARNING CONTEXTS

By

Ross Ian Walker

In order to keep pace with a rapidly changing environment, organizations must navigate a fundamental tension between exploration and exploitation. Over time, organizations often drift toward exploitation of known strengths and established resources, but this tendency can be harmful in a dynamic and competitive landscape. A classic simulation by James March (1991) demonstrated the importance of maintaining some degree of belief heterogeneity in an organization for the sake of long-term learning. In March's lineage, this thesis examines the effects of various exploratory strategies (i.e., individual experimentation, codification frequency, structural modularity, and employee turnover) on organizational learning in a bottom-up, networked, interpersonal learning context. Results demonstrate the complex interdependency of these variables in the exploration/exploitation tradeoff. Exploratory analyses suggest that a small degree of random individual experimentation has a favorable reward-to-risk ratio and that it is preferable to turnover as an exploratory strategy.

## ACKNOWLEDGEMENTS

I would like to thank my thesis chair, Dr. Richard DeShon, for his thoughtful guidance throughout this process. I would also like to thank the committee members, Dr. Steve Kozlowski and Dr. Kevin Ford, for their helpful comments.

## TABLE OF CONTENTS

|  |     |
|--|-----|
| LIST OF TABLES                                   | vi  |
| LIST OF FIGURES                                  | vii |
| INTRODUCTION                                     | 1   |
| Exploration & Exploitation                       | 4   |
| Conceptual Definitions                           | 5   |
| Exploitation Dominance                           | 5   |
| Cultural Transmission and Evolutionary Parallels | 7   |
| Background in I/O Psychology and Management      | 8   |
| A Typology of Strategy                           | 8   |
| Turnover as Variation                            | 9   |
| Organizational Learning                          | 11  |
| Modeling Organizational Learning                 | 11  |
| Why Use Computational Modeling?                  | 12  |
| March's Model                                    | 13  |
| Key Variables in the Model                       | 15  |
| Organizational Structure                         | 15  |
| Individual Experimentation                       | 19  |
| Tacit Knowledge                                  | 21  |
| Episodic Codification                            | 22  |
| METHOD   | 23  |
| Core Elements                                    | 23  |
| Individuals                                      | 23  |
| Environment                                      | 25  |
| Organizational Code                              | 25  |
| Organizational Structure                         | 25  |
| Knowledge  | 26  |
| Individual Knowledge                             | 27  |
| Organizational Knowledge                         | 27  |
| Knowledge Transmission                           | 28  |
| Learning From the Code                           | 28  |
| Learning By the Code                             | 29  |
| Interpersonal Learning                           | 29  |
| Experimentation                                  | 31  |
| Turnover & Environmental Turbulence              | 31  |
| RESULTS  | 32  |
| Tested Hypotheses                                | 32  |
| Exploratory Analyses                             | 34  |
| Experimentation vs. No Experimentation           | 35  |

|  |    |
|--|----|
| Experimentation vs. Turnover with Environmental Turbulence | 36 |
| DISCUSSION   | 37 |
| Tested Hypotheses  | 37 |
| Exploratory Simulations                                    | 39 |
| Limitations  | 41 |
| Future Research  | 43 |
| Practical Implications                                     | 45 |
| CONCLUSION   | 47 |
| APPENDICES   | 48 |
| APPENDIX A: Tables & Figures                               | 49 |
| APPENDIX B: R Simulation Code                              | 61 |
| APPENDIX C: Replication Discussion                         | 67 |
| REFERENCES   | 69 |

## LIST OF TABLES

|   |    |
|---|----|
| Table 1. Notional Representation of Initial Conditions                              | 51 |
| Table 2. Model Parameters and their Simulated Values in 3x3x3 Design                | 51 |
| Table 3. Effect Sizes Between All Conditions in the 3x3x3 Design                    | 53 |
| Table 4. Parameters in Exploratory Analyses and their Simulated Values              | 56 |
| Table 5. Ranges of $p_3$ in which Experimentation is Better than No Experimentation | 57 |

## LIST OF FIGURES

|  |    |
|--|----|
| Figure 1. Notional Representation of the Baseline Network Structure ( $\beta = 0$ )                            | 50 |
| Figure 2. Notional Representation of How Network Ties are Replaced   | 50 |
| Figure 3. Organizational Knowledge in the 3x3x3 Design   | 52 |
| Figure 4. Organizational Knowledge with No Turbulence and No Turnover  | 56 |
| Figure 5. Organizational Knowledge with Turbulence ( $p_{env} = 0.02$ ) and No Turnover                        | 57 |
| Figure 6. Organizational Knowledge with Both Turbulence ( $p_{env} = 0.02$ ) and Turnover ( $p_{turn} = 0.1$ ) | 58 |
| Figure 7. Replication of March (1991)  | 59 |
| Figure 8. Conceptual Replication of Fang et al. (2010)   | 60 |

## INTRODUCTION

Warren Bennis (1967) described bureaucracies as organizations with well-defined protocols and chains of command, specialized roles, an emphasis on technical competence, and an impersonal nature. He also predicted that bureaucracies would have died by now. Today, bureaucracy is alive but ailing (Davis, 2016). The murder is slow, as the four co-conspirators that Bennis named have been poisoning bureaucracy these last fifty years. Bennis' suspects include rapid change, globalization, the need for diverse skills, and humanistic management styles. Having an accurate model of reality becomes more challenging when rapid change creates a moving target. This study examines how levels of network connectivity, individual trial-and-error, and knowledge codification interact to influence organizational knowledge.

Traditional bureaucracies tend to perform best in stable, predictable industries that require efficiency through the management of clear, short-term problems (Burns & Stalker, 1961). More fluid, networked structures excel in dynamic and uncertain environments, as they can adapt with greater agility and follow through more effectively on long-term strategies (Kotter, 2012). More open communication and a greater propensity for risk-taking are conducive to organizational adaptation to change (Kontoghiorghes, Awbre, & Feurig, 2005). When what constitutes knowledge today could change tomorrow, organizations must be able to adjust strategies gracefully through continuous learning. Complex problems in a global economy will continue to need people who can pool their cognitive and social resources rather than go it alone.

Ashby and colleagues have highlighted the importance of an entity's degree of internal variety or degrees of freedom *in relation* to the environment. Applied to organizational science, the Law of Requisite Variety (Ashby, 1956) suggests that an organization must have the capacity



for at least as much variety as its environment does. That is, an organization's agility and internal complexity must equal or exceed the dynamism of the surrounding environment. Further, the Good Regulator Theorem (Conant & Ashby, 1970) suggests that maximizing both achievement and efficiency requires the development and maintenance of a sufficiently accurate model of reality. Therefore, a "well regulating" organization is aligned with its environment, not just through its structure but also through culture, beliefs, goals, etc.

When faced with rapid change and complexity, one of the greatest assets of a group is a diversity of perspectives. It seems intuitive that experts from different fields or different specializations in the same field should be able to collaborate to make better decisions than they could separately. However, research has found mixed effects of diversity in groups. Cohen and Levinthal (1990) describe the importance of balancing diversity and commonality in groups. They recognize that some degree of overlap in member knowledge is necessary to communicate effectively. For example, speaking different languages entails substantial process loss through translation or, potentially, an inability to communicate at all. Differences in values within a team can lower satisfaction and commitment (Jehn, Northcraft, & Neale, 1999). Certain types of demographic diversity can either provoke emotional conflict (Pelled, Eisenhardt, & Xin, 1999) or boost team satisfaction (Jehn et al., 1999). Functional background diversity can stimulate task conflict (Pelled et al., 1999) as can informational diversity (Jehn et al., 1999), and task conflict may drive team performance in these instances. Indeed, one meta-analysis showed only a very small, positive relationship between functional background diversity and team performance when examined across tasks (Bell, Villado, Lukasik, Belau, & Briggs, 2011).

The ambiguity in the literature on the effects of diversity in teams poses some important questions. Under what circumstances should a team or organization leverage the breadth of its

diversity to aim for optimal, long-term performance? Conversely, when should its members capitalize on the available knowledge for tangible, immediate results? These issues imply a fundamental tension that pervades individual, team, and organizational learning: the trade-off between *exploration* and *exploitation*. Performance of any type involves taking action based on current knowledge or the exploitation of existing resources. However, to succeed in a complex, interdependent task, a group often needs to leverage the diverse strengths of its members through collaboration and the cross-fertilization of ideas.

In a seminal paper, March (1991) computationally demonstrated the importance of balancing exploration and exploitation to optimize organizational learning. The simulations suggested that while maintaining diverse beliefs in an organization can slow progress initially, some dissent is vitally important for sustained organizational learning. Never acting on available knowledge produces no results, but having an especially strong organizational culture in which everyone quickly converges on the same “best practices” can yield suboptimal results.

The model in this current paper seeks a greater understanding of knowledge diversity in team-based organizational structures through the lens of the exploration/exploitation tradeoff. Encouraging diverse opinions and acting on new ideas is riskier than promoting homogeneity just as exploration is usually riskier than exploitation. Understanding how informational diversity manifests in organizations can encourage smarter risks in team and organization management. It can enable more effective balancing of short- and long-term perspectives and of micro and macro objectives. Many subtle mechanisms can drive organizations toward homogeneity, so this paper will focus most explicitly on examining the efficacy of various exploratory tactics. It does this with a computational model in March’s (1991) lineage and draws on notable extensions of March’s model (Fang, Lee, & Schilling, 2010; Miller, Zhao, & Calantone, 2006; Rodan, 2005) to

incorporate individual experimentation as well as interpersonal learning in a networked structure. As such, the study is well-situated in the organizational learning literature and focuses specifically on *intra*-organizational learning.

In summary, this computational study aims primarily to examine specific tactics that might balance exploration and exploitation across different conditions (e.g., internal structures, learning rates, and environmental change). Such tactics include individual belief change and the frequency of updating an organizational code. Previous models in this lineage have only examined learning from and by an organizational code or peer-to-peer learning – but never both. Since both mechanisms operate to some extent in almost all organizations, examining how they co-occur could productively guide theory building in the organizational learning literature. In addition to testing specific hypotheses on the balance between exploration and exploitation, this study also uses the model for more exploratory purposes that may inform future deductive research. After introducing foundational theory, offering hypotheses, and outlining the computational model, this thesis presents the simulated results and discusses their implications.

### **Exploration & Exploitation**

In organizational research, exploration and exploitation are somewhat nebulous concepts. However, they have proven useful in understanding adaptability (Mehlhorn et al., 2015). Most scholars conceptualize exploration and exploitation as two ends of a continuum implying that one should aim to balance the two (Gupta, Smith, & Shalley, 2006) rather than maximize orthogonal levels of both (e.g., Katila & Ahuja, 2002). Some call this coveted balance “organizational ambidexterity” (O’Reilly & Tushman, 2004). Additionally, classifying behavior as either one or the other can be deceptively difficult. A group decision to side with tradition might be exploitative in isolation but exploratory when contextualized by the group’s entire

lifecycle. Driven, conscientious employees wanting to learn the ropes in a new job are seemingly exploratory, but they could drive homogeneity and constitute an exploitative phenomenon at the organization level. Additionally, all activities deviate from prior ones in some way, even if imperceptibly, so some exploration is technically involved in even the most exploitative actions.

### **Conceptual Definitions**

Exploration is associated with search, novelty, learning, experimentation, and information gathering while exploitation is associated with execution, efficiency, and the use and refinement of existing resources or knowledge. Both of these are imperative to the adaptability and long-term survival of all species, as ignoring either one has distinctly negative consequences. Exploitation without exploration leads to the premature acceptance of mediocre or suboptimal results. Levinthal and March (1993) called this a “success trap.” Conversely, overvaluing exploration hinders any execution at all. This strategy has a much lower payoff expectancy and leads to a “failure trap.”

### **Exploitation Dominance**

While both strategies have advantages and disadvantages, March (1991) argued that exploration naturally carries greater risk and that the short-term attractiveness of exploitation can lure organizations toward conservatism. The payoffs from exploration are less certain and realized further into the future. Exploitation tends to produce clearer, more immediate feedback that can feed productive self-regulation. Specialized knowledge in one domain tends to increase the likelihood of attaining reward from its application. This, in turn, incentivizes further specialization. Investing in research and development does not guarantee a new revenue stream, but tightening an established manufacturing protocol to produce the same good at lower cost will almost certainly increase net income (holding all else constant). However, as March

demonstrated in his model, strict exploitation or incremental refinement of an existing resource can be maladaptive on a longer timeline and, therefore, myopic.

The attraction-selection-attrition cycle (Schneider, 1987) illustrates how homogeneity can emerge unintentionally. People are more interested in joining groups that share their goals, and groups tend to select people who share their goals. Even when divergent perspectives do enter the group, dissenters are more likely to leave due to irreconcilable differences. Each of these forces can restrict the range of knowledge, values, and beliefs in a team or organization. Schneider argued that this unintentional restriction of employee types can blind the organization to environmental changes or compromise its adaptability when changes are recognized.

Aligning teammates through shared mental models tends to improve team performance (DeChurch & Mesmer-Magnus, 2010), and this finding is somewhat intuitive. However, striving for a cohesive team culture in which everyone identifies strongly with a common mission could also stymie contradictory viewpoints that might have ultimately improved team effectiveness. Indeed, research suggests that minority dissent can improve the quality of group decisions and stimulate creativity (De Dreu & West, 2001). Sometimes combining people in teams yields worse outcomes than independent individuals, and a key factor in that phenomenon can be the group tendency to rush toward a suboptimal consensus without adequately considering divergent opinions (Hackman & Morris, 1975).

It is not always obvious when strategies and structures are encouraging too much homogeneity, traditionalism, and exploitation. Organizations are complex, dynamic systems involving a multitude of variables. Decisions rarely yield clear, immediate feedback, and the long-term consequences are generally far from intuitive. This is why exploitation of what we already know is attractive. It is also why exploration of what we do not know is an essential

counterbalance.

### **Cultural Transmission and Evolutionary Parallels**

Exploration and exploitation are analogous to variance and selection in evolutionary terms (Buss, 2012). Dawkins (2006) posited “When we die, there are two things we can leave behind us: genes and memes.” While gene frequencies drive human *biological* evolution, Dawkins argued that “memes” (e.g., ideas, songs, or in this case, beliefs) are the fundamental replicating unit in the human *cultural* evolution. Cultural and psychological variation may be more significant than genetic variation in modern times (Henrich & Boyd, 1998). Genes spread through sexual reproduction while memes spread between brains through social imitation. Genes that confer a greater survival advantage on their carriers enjoy greater prominence. Similarly, ideas and beliefs that bring an adaptive advantage tend to fare better along with their owners. Karl Popper (1972) advanced this evolutionary metaphor arguing that scientific knowledge advances through a process similar to Darwinian selection: hypotheses have different comparative fitness levels that influence their probability of survival.

March’s (1991) model and its extensions are evolutionary in nature such that more accurate beliefs tend to rise to the top through natural selection. On the other hand, some cognitive divergence is necessary to maintain the aforementioned requisite variety (Ashby, 1956). Fast learning from others can allow quick adaptation to immediate selection pressures, but it may hinder adaptation to future environmental shifts if it erodes internal variety. Fang et al. (2010) varied the complexity of problems that people faced such that an individual needed to hold a configuration of several correct beliefs for any of the correct beliefs to qualify as knowledge. This parameter is analogous to polygenic traits in which multiple genes determine phenotypic expression. The current model, however, assumes a simple, one-to-one

correspondence as with purely Mendelian traits.

### **Background in I/O Psychology and Management**

While the exploration/exploitation tradeoff is a widely applicable lens, it is rarely mentioned explicitly in organizational literature. Nonetheless, it has undergirded research and theory on key organizational phenomena including strategy formulation as well as turnover and selection.

#### **A Typology of Strategy**

Miles, Snow, Meyer, and Coleman (1978) recognized that administrative systems can have both lagging and leading functions. That is, organizations can crystallize useful information from the past and also create systems that allow innovation and future adaptation to occur. They proposed a typology of organizational strategy to make sense of the exploration/exploitation tradeoff. *Defenders* are exploitative. They are highly specialized organizations that carve out market niches and aggressively defend their turf against competitors. They do best in stable environments which reward efficiency and tend to have more rigid, top-down structures and highly divided labor. Conversely, the core strength of *Prospectors* is adaptability. They pursue novelty and innovation at the cost of efficiency and sometimes profitability. They allocate resources more disparately in experimental endeavors to stay relevant and lead the market. *Analyzers* represent a compromise between these two types and, potentially, the balance point between exploration and exploitation. *Analyzers* might maintain a core market and also adopt new technologies or products after *Prospectors* have proven their viability. Lastly, organizations that lack a clear strategy or stubbornly refuse to adapt to change are *Reactors*. Without a great deal of protection, *Reactor* organizations do not last long and must choose one of the first three strategies to survive.

Miles and colleagues (1978) argue that a “Human Resources” theory of management involving flatter organizational structures with more dispersed, bottom-up decision making is more viable in Prospector and Analyzer organizations. These types value adaptation and often cannot depend solely on a small group of leaders to succeed. Individual experimentation can relieve pressure on top management since it allows learning to occur at many levels of the organization (Hart, 1992). As environmental change quickens, organizations are forced toward more and more exploration to which traditional top-down management is poorly suited.

### **Turnover as Variation**

Typically, organizational scientists and HR practitioners see employee turnover as a problem, something to be minimized. Staw (1980) outlined the major downsides to turnover as (1) selection and recruitment, (2) training and development, (3) operational disruption, and (4) demoralization of organizational membership. Finding, vetting, and ultimately choosing new hires can be costly, and filling positions that require high levels of skill or experience often requires more time. Organizations with enough churn need to hire full-time staff just to carry out these functions. While the cost of a formal training program is self-evident, even informal peer training has an opportunity cost since it diverts time from mentors’ core job responsibilities. Additionally, it takes time for a new hire’s performance to reach an adequate and productive level and for the organization’s investment in the individual to realize a return. When an employee leaves, his or her extant colleagues’ productivity may suffer, particularly when the work is highly interdependent and when the vacant role is highly specialized or high status in the organization. Lastly, turnover can have attitudinal costs for remaining employees if they perceive the exit as a reflection of systemic problems in the organization. In this way, turnover can potentially trigger more turnover.



However, Staw (1980) also highlighted the underemphasized benefits of organizational turnover: (1) increased performance, (2) reduction of entrenched conflict, (3) increased mobility and morale, (4) innovation and adaptation. Sometimes new hires perform better than their predecessors, and young employees can bring a renewed motivational force. Turnover can reduce interpersonal clashes that are either unresolvable or not worth the effort to resolve. Turnover can also have positive attitudinal consequences when employees notice their colleagues get better opportunities, as it signals the value of their current employment for future goals. Most germane to this paper, when turnover leads to hiring outside the organization, new people can bring new ideas, skills, and experiences. While it is certainly not guaranteed, this injection of variety can support innovation and more effective adaptation to a changing environment by subjecting entrenched beliefs, values, and operating procedures to “naïve” scrutiny. These positive consequences of turnover are harder to quantify, and they are realized less immediately than the downsides. Nonetheless, they can still contribute meaningfully to the long-term viability of an organization.

While maintaining the same employees and the same conception of reality that has always worked is the optimal strategy in a static environment, March (1991) demonstrated computationally the positive effects of turnover in a dynamic environment. When reality changes, previously correct beliefs grow stale, so introducing new agents with randomly assigned beliefs allowed the organization’s codified knowledge level to maintain a relatively high level rather than sinking into maladaptation. Similarly, March’s model showed that turnover can counteract the homogenizing effect of high socialization rates (i.e., when employees are too quick to adopt the organizational prescribed beliefs). The current model will examine both turnover and individual belief change as exploratory mechanisms which introduce stochastic

variation into an organizational system.

### **Organizational Learning**

Despite a lack of clarity on exactly how to define it, the topic of organizational learning has supported interesting theory, research, and discussion. Psychologists have typically studied the process of organizational learning while management researchers have focused on the outcome of that learning (i.e., sustainable competitive advantage; Dodgson, 1993). In line with Ashby's position, organizations must develop and maintain a sufficiently accurate model of a dynamic reality to adapt and survive. Learning is typically conceived as an individual activity, so individual learning is the fundamental substrate of organizational learning. However, as Hedberg described, "Organizations do not have brains, but they have cognitive systems and memories...Members come and go, and leadership changes, but organizations' memories preserve certain behaviours, mental maps, norms, and values over time" (1981, p. 3). Argyris and Schon (1978) argued that, at bottom, an organization is defined by its shared goals, so the state of such organizational characteristics which dictate daily functioning are critical. Organizational learning scholars have also emphasized the importance of context (e.g., what people learn depends on the people and the environment that surrounds them; Simon, 1991) and time (e.g., discarding obsolete information when appropriate is just as important as acquiring new information; Hedberg, 1981).

### **Modeling Organizational Learning**

Senge (1995) describes learning organizations as continually and collectively working toward greater understanding of relevant systems and greater capacity for sustainable success. However, learning as a psychological process is difficult to capture at the individual level and even more daunting at the team and organization levels. Indeed, researchers have often treated

team learning as an unknowable “black box” (Grand, Braun, Kuljanin, Kozlowski, & Chao, 2016). Some have used team performance as a proxy for learning (Kozlowski & Ilgen, 2006) while others operationalize team knowledge as the sum or average of individual knowledge (Bell & Kozlowski, 2008). This suggests a conceptualization of team learning as a precursor to emergent, compositional outcomes like team mental models (Cannon-Bowers, Salas, Converse, 1993). This aligns with what DeChurch and Mesmur-Magnus (2010) call “cognitive similarity-congruence” or the degree of convergence of member cognitions. However, this conception of team learning does not reflect its true complexity, and computational modeling may be the only method currently available to study such complexity. This complexity is only magnified in organizational learning, and it explains why March (1991) conducted his study “in silica” rather than in the field.

### **Why Use Computational Modeling?**

Computer simulation is still an underutilized research paradigm in organizational science (Weinhardt & Vancouver, 2012) in which computational models of systems are paired with experimental designs (Harrison, Lin, Carroll, & Carley, 2007). This approach has many advantages that supplement traditional research methods. In this study, the two most salient are that simulations can (1) examine how complex systems with many variables function over time, and (2) allow and encourage development of clearer, more internally consistent theory.

Accurately assessing the complex dynamics of teams and organizations is often intractable in empirical studies. Just thinking through the dynamic interdependencies of organizational subsystems quickly runs into a processing ceiling of the human mind. In terms of raw processing power, computers are far more robust, and that capacity is growing rapidly. Further, the relatively few studies that do measure change over time often rely on arbitrary

sampling frequencies or ones based on convenience (Hulin & Ilgen, 2000). Computer simulations can help make sense of complex phenomena, but few researchers have computationally modeled the exploration/exploitation tradeoff with realistic representations of organizational structures. The empirical studies that have examined it at the team level may have limited generalizability (e.g., Taylor & Greve, 2006; Peretti & Negro, 2006). The potential for computational modeling to reveal non-obvious, emergent phenomena from systems in which the data generating mechanisms are fully specified can guide theoretical progress in organizational learning (Miner & Mezias, 1996).

Informal theories conveyed in natural language usually have ambiguities that formalized or mathematically represented theories do not (Vancouver & Weinhardt, 2012). Ilgen and Hulin (2000) identify two types of “overidentification errors.” The first occurs when a multitude of different theories seem to explain the same phenomena. In this case, the observations have too little information to differentiate good claims from bad ones. The aforementioned advantage of computational modeling addresses this potential error by providing essentially perfect observational fidelity. The second overidentification error occurs when a theory is not adequately falsifiable. Here, it is the theory that has too little information. Its ambiguity maintains immunity from contradiction. Formalizing theory in equations or concrete rules and testing them computationally guards against this potential error. This strategy “quickly reveals logical inconsistencies that lie hidden beneath the verbal turf” (Hulin & Ilgen, 2000).

### **March’s Model**

March (1991) used an elegant computational model to show that slow socialization to the organization’s code of knowledge combined with a code that adapts quickly to reflect the best available information yields the best long-term learning outcomes. While knowledgeable

employees can influence the organizational code in this scenario, all learning happens through the code and is, therefore, entirely top-down. When individuals learn slowly from the organization, greater cognitive diversity persists which allows the code to eventually become more accurate or more concordant with reality than if every newcomer socialized quickly and thoroughly. Preserving nonconforming beliefs through slow learning at any organizational level is exploratory, as it forsakes existing knowledge to prospect for a better future. Fast learning from the organizational code can increase aggregate knowledge quickly, but it sets a lower ceiling on what the group can ultimately achieve.

March's (1991) model begins with a set of randomly determined "truths" about reality. Agents in the model begin with randomly assigned beliefs about each of these dimensions. Agents with the most accurate beliefs are assumed to be superior performers, so they ascend to the policy-making elite that dictates an organizational code. Specifically, anyone whose knowledge exceeds that of the code is allowed to influence it. The code structure is isomorphic to reality and the beliefs of individuals but represents formally communicated organizational knowledge or best practices. The model includes two learning mechanisms with associated probabilities: individuals learning from the code and the code learning from the policy-making elite. While this bidirectional learning process can fuel long-term success, an imbalance can easily impede it. The primary marker of success in this model is organizational knowledge or the degree of concordance between the organizational code and reality.

March (1991) included two parameters that would introduce stochastic shocks in an open system: turnover and environmental turbulence. Exogenous probabilities determined when a new hire with random beliefs ousted a current employee and when the dimensions comprising reality changed. Argote, Insko, Yovetich, and Romero (1995) noted the adverse effects of turnover on

team knowledge. However, while the replacement of a veteran with an outsider who is unfamiliar with team processes and shared knowledge can certainly drain what Simon (1991) called “organizational memory,” turnover can also introduce valuable variance. Indeed, March found that low to moderate rates of turnover could ameliorate the homogenizing effect of fast learning from the organizational code and maintain higher aggregate knowledge levels in a turbulent environment. He cautioned that “A major threat to the effectiveness of [mutual learning] is the possibility that individuals will adjust to an organizational code before the code can learn from them.”

### **Key Variables in the Model**

#### **Organizational Structure**

Geneticist Sewall Wright (1932) noted that evolving populations typically show loose coupling or modularity. He explains that in nearly isolated subgroups that allow some crossbreeding, “all gene frequencies can drift irregularly...without reaching fixation and giving the effects of close inbreeding. The resultant differentiation...is of course increased by any local differences in the conditions of selection.” The persistence of relatively small, isolated subgroups maintains genetic diversity that would diminish if everyone was part of one large, fully connected population. In this latter case, Wright says, “further evolution can only occur by the appearance of wholly new (instead of recurrent) mutations...which happen to be favorable from the first.” In this scenario, full interconnection with low variability produces a higher average but a likely forfeiture of optimality. A study in computational evolution found that networks performed better when incentivized to minimize connection costs between nodes (Clune, Mouret, & Lipson, 2013). The additional constraint led to more modular network structures that were more effective and “evolvable.” The authors listed potential reasons for this phenomenon

including fewer parameters to optimize and faster, more sustainable adaptation since negative shocks could stay confined in individual subsystems without spreading to the entire population. Additionally, Lipson (2007) offered the caveat that “increased performance gained by reduction of modularity is often justified in the short term, whereas increased modularity is often justified over longer time scales where adaptation becomes a dominant consideration.” This aligns with the established notion that exploitative strategies often perform best in the short term by rising quickly to equilibrium, but exploratory strategies can achieve a higher equilibrium and greater performance in the long term.

Loose connectivity has many advantages in organizations. Scholars (e.g., Weick, 1976; O’Reilly & Tushman, 2004) have argued that structural separation can help maintain adequate cognitive diversity and balance exploration and exploitation. For example, Benner and Tushman (2003) found that isolating new product design teams from the strictures of organizational norms helped innovation flourish through the exploration of new alternatives. Weick (1976) explained that “It is conceivable that loosely coupled systems preserve more diversity in responding than do highly coupled systems, and therefore can adapt to a considerably wider range of changes in the environment...” Fiol and Lyles (1985) summarize this point well: “A centralized, mechanistic structure tends to reinforce past behaviors, whereas an organic, more decentralized structure tends to allow shifts of beliefs and actions.” Additionally, modularity in a system can limit the potential damage that unforeseen shocks can bring (Weick, 1976; Page-Jones, 1980). Herbert Simon’s (1962) parable of the watchmakers demonstrates this benefit of modularity. The watchmaker whose method includes separable, independent chunks does not suffer disturbances as much as the peer who needs to restart each watch from scratch every time the phone rings. Also, given a fixed set of inputs (e.g., employees), modularity allows a higher number of

possible configurations which can increase the flexibility of the whole system (Schilling, 2000).

Loose coupling also has weaknesses. Just as problems are less likely to spread between units, useful ideas or positive mutations are also less likely to spread (Weick, 1976).

Additionally, while many systems tend to converge toward loose coupling, this is not universally true. Schilling (2000) outlines a theory that explains how and under which circumstances systems converge either toward or away from modularity. When the inputs to and demands of a system are relatively homogenous, modularity can degrade fitness.

While most empirical research has examined subgroups with exploratory missions, Miller et al. (2006) showed computationally that physical separation and local learning could preserve diversity while unfettered distant learning suppressed it. Their model extended March's (1991) in several ways including the introduction of direct, interpersonal learning and tacit knowledge transfer independent of the organizational code. Agents could learn locally from agents in direct physical proximity. However, when agents decided they had nothing left to learn from their immediate teammates, they could also explore the organizational network more broadly. Through distant learning, agents gained access to the whole organizational network. This strongly homogenized the organization, as unrestricted distant learning rendered learning from the organizational code obsolete.

Fang et al. (2010) demonstrated how semi-isolated subgroups can maintain requisite variety while selecting for greater knowledge over the long-term. Their model includes direct, interpersonal learning from team members and between semi-randomly linked teams. They used a slightly modified version of the "connected caveman" model (Watts, 2009) in which otherwise isolated cliques or "caves" have a very small degree of connectivity between them in the baseline model. The point of optimal learning that Fang et al. (2010) found corresponds theoretically to



what Wright (1932) described as being most conducive to long-term evolutionary adaptation. Kotter (2012) argues that organizations can benefit from having two complementary structures operating simultaneously. One is a traditional hierarchy that handles concrete, short-term tasks and the other is a “voluntary army” of people from different subgroups and levels working on strategy and long-term adaptability. From this lens, the  $\beta$  parameter that Fang et al. (2010) used to add inter-team ties to their baseline network is analogous to Kotter’s strategic network.

Fang et al. (2010) found that no connectivity beyond the basic structure yielded the lowest aggregate knowledge followed by highly connected subgroups and then by semi-isolated subgroups (i.e., when 10% of *intra*-team ties were repurposed as *inter*-team ties). This preserved subgroup identity through adequate modularity but still allowed knowledge to spread to other teams. They found that this loosely coupled subgroup structure was ideal for aggregate learning across a range of contingency variables. Therefore, they interpreted their findings as evidence against contingency theorists (e.g., Burns & Stalker, 1961) who would argue that no ideal organizational structure exists since it depends on a variety of factors like organizational goals and environmental stability. Contingency theorists might also argue that one structure can be exploratory or exploitative depending on other moderating factors. The findings of Fang et al. (2010) do not rule out contingency theory since it is possible that this pattern emerged as an artifact of the learning mechanism’s operationalization. However, this model will attempt to replicate conceptually this ideal point of modularity.

*Hypothesis 1: With no organizational code and no individual experimentation, a small degree of inter-team connectivity (i.e., “loosely coupled subgroups” with 10% of intra-team ties repurposed as inter-team ties) will yield higher levels of organizational knowledge than isolated and tightly coupled subgroups.*

## **Individual Experimentation**

When subgroup isolation fizzles (i.e., the organization becomes increasingly interconnected), long-term knowledge tends to decline slightly as the organization shifts toward exploitation. To mitigate this driver of homogeneity, many models inject variance through employee turnover and environmental turbulence. The current model will include these but also focus on individual experimentation as another potentially useful source of variation.

Most extensions of March's (1991) model assume that all learning in an organization occurs either interpersonally or from formalized knowledge at the organizational level. Huber (1991) noted that almost no work in organizational learning had focused on unintentional or unsystematic learning. People can acquire knowledge anywhere, and the augmentation of one's beliefs need not be constrained to working hours. Rodan (2005) incorporated various forms of experimentation but found that only random "foolishness" had a significant and positive effect on learning outcomes. Another form was organizationally constrained experimentation which only allowed mutations to beliefs on which the organizational code was neutral. Lastly, self-restrained experimentation only allowed mutations to beliefs on which the individual was neutral. The blind variation from completely random experimentation outperformed both of these more deliberate and thoughtful strategies. Like Rodan's, this model will allow individuals to change their beliefs irrespective of the organizational code, interpersonal interaction, or preexisting belief structures. For this reason, the "experimentation" in this model has no pre-specified goal, nor does it intentionally seek feedback on whether the change was effective. It simply represents individuals changing their beliefs, maybe as a consequence of personal experience or research. While this random experimentation is not an intentional learning process, the organization as a whole may be able to learn from such "foolishness."

In reinforcement learning, the multi-armed bandit problem sheds light on the exploration/exploitation dichotomy (Sutton & Barto, 2012). Slot machines, also known as “one-armed bandits,” produce some payoff at a particular probability. The multi-armed bandit problem tasks an agent with using multiple metaphorical slot machines with varying payout probabilities to maximize total winnings over a limited number of trials. The agent does not know the expected value of any machine’s payout but may develop estimates through experimentation. The action that the agent estimates to have the highest value at a given time is called the *greedy* option. To maximize payoff on the next pull, a rational agent chooses the greedy option to exploit its current knowledge. On the other hand, the agent might choose to explore other options and potentially sacrifice short-term payoff to find a machine with a greater expected value. How should an agent approach this decision? One solution is to exploit the greedy option most of the time but experiment with a small probability  $\epsilon$ . This strategy is known as  $\epsilon$ -*greedy*. As the aforementioned models would predict, no exploration (i.e.,  $\epsilon = 0$ ) tends to show faster initial improvement, but slight experimentation yields greater long-term winnings.

This example demonstrates how introducing random variation in action selection problems can lead to more desirable results. In this proposed model, individuals will alter their beliefs at some specified rate to maintain variation in their core team and in the system overall. In evolutionary terms, individual experimentation is akin to random genetic mutation. Some (e.g., Eiben & Schipper, 1998; Črepinšek, Liu, & Mernik, 2013) have argued that mutation in evolutionary algorithms can be exploitative because most of the prior material remains. However, it is still inherently exploratory, as it introduces unbiased novelty. In tightly coupled subgroups that promote fast convergence toward a homogenous, suboptimal organizational knowledge level, individual experimentation should counteract the homogeneity through random

belief mutations. Conversely, nearly isolated subgroups maintain dissent well. Rather, these groups can struggle to capitalize on superior knowledge in other parts of the organization. It should follow that individual experimentation would exacerbate their isolation and that low levels of experimentation should produce better outcomes.

*Hypothesis 2a: In tightly coupled subgroups, high individual experimentation will yield higher organizational knowledge than low experimentation and no experimentation.*

*Hypothesis 2b: In disconnected subgroups (i.e., the baseline structure), no individual experimentation will yield higher organizational knowledge than low experimentation and high experimentation.*

### **Tacit Knowledge**

Bandura's (1963) social learning theory argued that we learn not just through direct instruction but also through observation and mere proximity. Ostroff & Kozlowski (1992) noted that newcomers socialize primarily through informal observation of others and experimentation. Polanyi (1967) posited two categories of knowledge: explicit and tacit. Explicit knowledge can be expressed directly in a formal language while tacit knowledge is transmitted only through indirect exposure and experience. The latter forms through implicit learning or "the process through which one becomes sensitive to certain regularities in the environment: (1) without trying to learn regularities, (2) without knowing that one is learning regularities, and (3) in such a way that the resulting knowledge is unconscious" (Cleeremans & Dienes, 2008). While most education focuses on the formalized transmission of directly communicable knowledge, it seems that a significant proportion of learning in organizations happens in subtler ways.

March's (1991) model assumes explicit learning through a formalized organizational code, but Miller et al. (2006) allowed tacit dimensions by restricting the number of beliefs that

the code included. These inert dimensions of the code meant that their knowledge would need to spread interpersonally rather than top-down. This proposed model will adopt this assumption since it is more realistic, and it weighs interpersonal learning more heavily.

### **Episodic Codification**

Beyond the code's inability to transmit all useful knowledge, March, Shulz, and Zhou (2000) note that, in reality, code updates occur periodically, not continuously. March's (1991) original model assumed constant updating of the organizational code according to the beliefs of the policy-making elite and constant learning from the code by employees. Miller et al. (2006) addressed this limitation by including episodic updates and limiting the timeframe during which agents could learn from it. When the code updated, people learned what they would based on the learning rate and then ignored the code until its next iteration. This proposed model will adopt a similar codification scheme. Both of these features, along with the allowance for tacit knowledge transmission, lessen the code's influence and tend to preserve heterogeneity. This exploratory effect is expected to interact with organizational structure.

*Hypothesis 3a: In disconnected subgroups, frequent code updates will yield higher organizational knowledge than infrequent and non-existent code updates.*

*Hypothesis 3b: In tightly coupled subgroups, non-existent code updates will yield higher organizational knowledge than frequent and infrequent code updates.*

## METHOD

The current model incorporates elements from March's (1991) original model and from the three aforementioned extensions (Fang et al., 2010; Miller et al., 2006; Rodan, 2005). Broadly, the model focuses on interpersonal learning, a simplified organizational code that allows for tacit knowledge accrual, and individual experimentation or learning from outside of the organization. To test the proposed hypotheses, simulations mimic a 3x3x3 experimental design by varying organizational structure, individual experimentation, and the frequency of code updates. In addition to testing the stated hypotheses, this paper also explores the emergent phenomena in the model inductively to decipher interesting patterns that might inform future theory. The model was constructed in the open source statistical program *R*. Broadly, a script establishes initial conditions and runs functions aligned with the equations outlined in the following sub-sections. Those functions generate and store longitudinal data in matrices for analysis. The full script is included in Appendix B.

### **Core Elements**

Faithful to March, the model has three core elements: individuals, an exogenous reality (also known as “the environment”), and an organizational code.

#### **Individuals**

Dodgson (1993) said that “individuals are the primary learning entity in firms, and it is individuals which create organizational forms that enable learning in ways which facilitate organizational transformation.” The model includes  $n$  individuals organized into teams of  $z$  individuals each. Each individual holds  $m$  beliefs and begins the simulation with each belief set randomly to 1, 0, or -1 (i.e.,  $b^i = b^i_1, b^i_2, \dots, b^i_m$  with  $b^i_j \in \{-1, 0, +1\}$ ). Holding a belief of

$b^i_j = 0$  indicates neutrality or indecision on that dimension.

Fang et al. (2010) found that at the optimal level of inter-team coupling (i.e.,  $\beta = 0.1$ ), seven-member groups yielded the optimal solution, so individuals will be organized into teams of seven by default. Sensitivity analyses revealed that lower team sizes (e.g., teams of two agents each) produced slightly lower aggregate knowledge levels. However, the results from seven-person teams used here produce results that are representative of a wider range of larger team sizes. This also seems to be a reasonable team size given prior research in both traditional and modern organizations. Bantel and Jackson (1989) found an average team size of 6.3 people for top management teams in banking. More recently, archival data estimated the average team size in software development to be 7.9 people per team (Rodriguez, Sicilia, Garcia, & Harrison, 2011).

Fang et al. (2010) ran simulations with 100 belief dimensions and 200 runs per condition, and this model does the same. Simulations with more beliefs yielded less variable aggregate knowledge levels, but they also yielded lower aggregate knowledge, presumably because reality is more complex and, therefore, more difficult to know. Running 200 trials per condition ensures adequate stability of knowledge outcomes. Also, including 200 time points per simulation allows organizational knowledge to reach a more satisfactory equilibrium (e.g., rather than ending a trial in the middle of an ascent or descent). Fang et al. (2010) began with 280 individuals in the model, but this model begins with 140 individuals to represent a medium sized organization. This smaller organizational size is still relatively large given that 89.24% of organizations in the U.S. have fewer than 20 people and 98.15% have fewer than 100 employees (United States Census Bureau, 2015). However, an organizational size of 140 produced a more stable estimate of other larger organizational sizes whereas a smaller organizational size of 70 agents produced

slightly higher levels of average individual knowledge. The number of teams in the organization is represented as  $n / z$ .

### **Environment**

Reality also has  $m$  dimensions that are randomly determined but do not include zero (i.e.  $e = e_1, e_2, \dots, e_m$  with  $e_j \in \{-1, +1\}$ ).

### **Organizational Code**

This code is a belief set from which agents can learn. It has the same structure as the belief sets of individuals and the environment. This code begins all simulations with all dimensions set to zero. However, in some instances, the code only includes a subset of the  $m$  dimensions that the environment and individuals have. The parameter  $q$  dictates the number of tacit dimensions in the code. These dimensions remain at zero through the simulations and are, therefore, non-functional. That is, the organizational code can be characterized functionally as  $c = c_1, c_2, \dots, c_{m-q}$  and  $c_j \in \{-1, 0, 1\}$  since all other dimensions of the code will be neutral (i.e.,  $[c_{m-q}, c_m] = 0$ ).

Miller et al. (2006) specified that half of all beliefs should be tacit. However, unsurprisingly, when learning was more tacit, the frequency of organizational code updates had less of an effect on knowledge. The current model uses  $q = 50$  as its default.

### **Organizational Structure**

In the baseline structure, each team is fully connected within itself (i.e., each team begins with  $\frac{z(z-1)}{2}$  total ties) but fully disconnected from other teams. For every tie that a focal individual has to a teammate, the  $\beta$  parameter represents the probability that that intra-team tie is eliminated and replaced with a tie to an individual outside of the team. This process repeats for every tie and for every individual in the organization sequentially. The result is that high levels



of  $\beta$  indicate greater *inter*-team connection, less *intra*-team connection, and therefore, less overall clustering. As the protocol behind the  $\beta$  parameter suggests, the total number of ties in the network does not change from the baseline model. Each new tie is preceded by the elimination of one. Figure 1 is a notional representation of the baseline network structure, and Figure 2 exemplifies how ties are replaced randomly according to  $\beta$ . Structure is determined before each trial, but it does not change over the course of a trial (i.e., individuals cannot form or drop social ties). The only difference between this protocol and the one used by Fang et al. (2010) is that the baseline model (i.e.,  $\beta = 0$ ) here has completely disconnected teams while the baseline, connected caveman structure began with minimal ties between neighboring teams. Beginning with isolated teams makes the  $\beta = 0$  instance more meaningful.

The model tests three structures with values of  $\beta$  equal to 0, 0.1, and 0.5. Respectively, these values correspond to the disconnected, loosely coupled, and tightly coupled subgroup structures in the hypotheses. For each run, the simulation generates a binary, symmetric matrix with  $n$  rows and  $n$  columns. Because relationships in the model are always bidirectional and do not vary in strength (neither between relationships nor between directions in a relationship), a 1 in the matrix indicates that the two members are linked, and a zero indicates that they are not linked.

### **Knowledge**

Knowledge is operationalized as the concordance of a belief set with the environment. To determine whether or not a belief is correct, the coding allows simple multiplication of the individual belief with the corresponding dimension of reality (i.e.,  $b^i_j e_j$ ). This formulation yields -1 when an individual chooses the wrong belief (e.g.,  $b^i_j = 1$  when  $e_j = -1$ ), 0 when the individual is neutral, and 1 when the individual's belief is concordant with the environment. In

addition to holding a belief that is true, Baron (2000) asserts that true knowledge also requires the holder to have chosen the belief based on the right evidence and inferences. That is, the belief should be both true and justified. This model ignores how and why individuals choose beliefs, however, and concordance with reality is the sole criterion.

### **Individual Knowledge**

This is the degree of concordance between an individual's belief set and reality. Averaging the products of an individual's beliefs and their corresponding environmental dimensions yields individual knowledge such that  $IndKnow^i = \frac{1}{m} \sum_{j=1}^m b^i_j e_j$ . It follows then that  $IndKnow^i$  and all other knowledge levels are bounded by -1 and 1, and the expected value  $E(IndKnow^i) = 0$  at time  $T = 0$ .

### **Organizational Knowledge**

Aggregate organizational knowledge across conditions will be the primary outcome of interest. To calculate this, two levels of intermediary averages are required:

1) *Time*: the average of every individual's knowledge at each time point in each trial.

Formally, this is calculated as  $AvIndKnow = \frac{1}{n*m} \sum_{i=1}^n \sum_{j=1}^m b^i_j e_j$ .

2) *Condition*: the *Time* averages are themselves averaged across  $N$  trials per condition (this produces the typical trajectory over time for each condition).

Finally, the aggregate organizational knowledge level for a condition is the average of the knowledge levels at each time point. This metric is then divided by  $m$  (i.e., the number of dimensions in reality) to yield the proportion of possible knowledge. While this value can technically range from -1 to 1, in practice, it only dips barely below zero at  $T = 0$  due to the random configuration of beliefs in the initial conditions.

March (1991) operationalized knowledge as the concordance of *the code* with reality.

While having leadership with high knowledge levels would be likely to improve organizational performance, it is also possible for management to have all the answers while the employees ignore the knowledge. This model chooses to operationalize organizational knowledge differently for two reasons. First, the organizational code is meant to be a feature or only one possible way to learn in this more bottom-up learning context. Code knowledge as an outcome is incongruent with that goal of the model, and examining the aggregate of all individual knowledge levels is a more appropriate measure. March and Fang et al. (2010) both used equilibria as an outcome measure in various analyses (i.e., the level at which code knowledge leveled out and stayed roughly constant). This model eschews measuring equilibria because it carries an unfair bias toward exploratory strategies. In a static environment, exploitative strategies tend to rise quickly to a mediocre equilibrium. A more long-term, exploratory strategy takes longer to rise knowledge levels but often reaches a higher equilibrium. Simply measuring the end state of these two trajectories ignores the slower rise of exploration and the opportunity cost that it carries. Thus, organizational knowledge in this model accounts for knowledge levels throughout the course of each run.

### **Knowledge Transmission**

There are three main learning mechanisms in the model: learning *from* the code, learning *by* the code, and learning from others (interpersonal learning).

#### **Learning From the Code**

This model uses March's (1991) parameter,  $p_1$ , to indicate the probability in any given time period that an individual adopts a belief from the organizational code. This is set exogenously and applied uniformly to all individuals.

## Learning By the Code

The code reflects the majority view of a policy-making elite. This group consists of individuals with a knowledge level that exceeds that of the organizational code.  $k^c_j$  reflects the number of policy makers who agree on the  $j^{th}$  dimension of the code subtracted from the number of policy makers who disagree. Therefore, positive values of  $k$  indicate that the code will change at some probability. However, when  $k$  is negative (i.e., more people agree than disagree) or zero (i.e., there is a tie), the code does not change. The value of  $k^c_j$  is used to determine  $v^c_j$  (the position of the organizational code on dimension  $j$ ), and this transformation is formally expressed as:

$$v^c_j \in \{-1, 0, +1\} = \begin{cases} +1 & \text{if } k^c_j > 0 \\ 0 & \text{if } k^c_j = 0 \\ -1 & \text{if } k^c_j < 0 \end{cases}$$

The probability of  $c_j$  changing to  $v^c_j$  is  $1 - (1 - p_2)^{k_j}$ . Table 1 depicts a simplified and hypothetical scenario with randomly generated starting values in R. The organizational code begins completely neutral, so any individual with a positive knowledge level will participate in the policy-making elite until the organizational code surpasses the individual. Additionally, tacit dimensions of the code remain at zero throughout the trials.

Whenever the code updates, people have the opportunity to learn from it. The  $\tau$  parameter represents the frequency of code updates. Therefore, the tested values of 0, 1, and 10 represent no code, continuous updating, and sporadic updating once every 10 time points respectively. These values allow exploration of different orders of magnitude while controlling the number of conditions.

## Interpersonal Learning

Individuals can learn from the people to whom they are connected, whether in their core

team or not. Individuals can decipher the aggregate knowledge level of a peer, but they do not know which of the peer's beliefs are correct and which are not. The interpersonal learning mechanism in this model reflects both a prestige bias and a conformist bias (Henrich, 2001). Prestige bias suggests that people tend to mimic the beliefs of superior performers, not because the beliefs are better but because the people have higher status. A conformist bias exists when people favor ideas espoused by a majority of others over those expressed by a minority (e.g., Davis, Kerr, Atkin, Holt, & Meek, 1975). Therefore, individuals will learn only from people who are linked to them in the network and have higher knowledge levels than themselves. Further, individuals will only ever adopt the majority view of their superior others. With this learning mechanism, groups tend to converge toward homogeneity at first. Once an individual knows as much or more than others in the group, the influence of out-group members tends to grow. The choice to operationalize interpersonal learning through a majority of superior peers rather than only from the highest performing peer highlights an important assumption. Fang et al. (2010) note that the latter strategy significantly attenuates the effect of subgroup structure on performance. In fact, such a "tournament selection" rule tends to produce the opposite results: higher performance at very low and very high levels of subgroup connectivity. They operationalized interpersonal learning in this way, and the current model attempts to replicate it.

The "majority rule" decision-making process requires several pieces of information. First, each dimension in each agent is either selected for learning or not according to the parameter  $p_3$ . This parameter is similar to the coefficient of imitation in a Bass diffusion model which is typically set between 0.3 and 0.5 with a mean value of 0.38 (Mahajan, Muller, & Bass, 1995). Fang et al. (2010) used a default of  $p_{learning} = 0.3$ , and this model will do the same for  $p_3$ . Agents then recognize who among their peers has superior knowledge to themselves (i.e. those

peers with  $IndKnow^i$  exceeding their own). As with the organizational coding process, the individual determines the majority position on each dimension  $j$  within the higher performing peers by summing the values of each superior team member to obtain  $k_j^l$ . Values of  $k_j^l$  determine  $v_j^l$ , or the majority's view on dimension  $j$  such that:

$$v_j^l \in \{-1, 0, +1\} = \begin{cases} +1 & \text{if } k_j^l > 0 \\ 0 & \text{if } k_j^l = 0 \\ -1 & \text{if } k_j^l < 0 \end{cases}$$

If a dimension was selected for learning and the associated agent has peers with superior knowledge levels, the agent will adopt  $v_j^l$ . Table 2 shows planned starting values for simulations that will test the hypotheses.

### **Experimentation**

The probability that any belief held by any individual will change is determined by the parameter  $\varepsilon$ . A belief selected to change will always change to either a 1 or a -1. If the belief was originally set to 0, then either 1 or -1 will be selected at equal (i.e., 0.5) probabilities. The model will explore  $\varepsilon$  levels of 0, 0.01, and 0.1 to examine different orders of magnitude.

### **Turnover & Environmental Turbulence**

Turnover and changes to the dimensions of reality follow the same protocol as March's (1991) model. The parameter  $p_{turn}$  represents the probability that any given individual in the model is eliminated and replaced with a "naïve" individual with randomly set beliefs. For each dimension of reality,  $p_{env}$  represents the probability that it will change at each time point (i.e., from a 1 to a -1 or vice versa). A  $p_{env}$  value of 0.02 suggests that, on average, two dimensions out of 100 will change at each time point. Therefore, with  $T = 100$  time points, the environmental code would flip entirely two times over the course of a trial—*on average*.

## RESULTS

### Tested Hypotheses

To test the hypotheses, a 3x3x3 experimental design was used to manipulate the connectivity of teams ( $\beta$ ), individual experimentation ( $\varepsilon$ ), and the frequency of code updates ( $\tau$ ). In each condition, the focal criterion is organizational knowledge operationalized as the average concordance of employee beliefs with reality across each time point and each trial. Because some mechanisms of this model are based on predecessors, two replication experiments were conducted to replicate key results from March (1991) and Fang et al. (2010). The full discussion of those replications can be found in Appendix C.

Figure 3 shows aggregated organizational knowledge across the three levels of inter-team connectivity, code update frequency, and individual experimentation. When there is no code ( $\tau = 0$ ) and no experimentation ( $\varepsilon = 0$ ), loosely coupled subgroups ( $\beta = 0.1$ ) outperform isolated subgroups substantially and tightly coupled subgroups marginally (Table 3 includes all effect sizes for context). This degree of clustering also produced the highest knowledge level when the code updates every 10 time points. However, when the code updates at every time point, there is no significant difference between loosely coupled and tightly coupled subgroups. Therefore, hypothesis 1 is supported, but loosely coupled subgroups do not outperform tightly coupled ones across all conditions. Even with the slightly different network structure (i.e., no inter-team ties at  $\beta = 0$ ), these results replicate the superiority of loosely coupled subgroups found in Fang et al. (2010).

Hypothesis 2a is partially supported, since high individual experimentation ( $\varepsilon = 0.1$ ) yields the highest organizational knowledge level in tightly coupled subgroups when the code

updates at every time point. However, as shown in the left and right panes of Figure 3, high experimentation has a highly negative effect on learning when the code is less active. A constantly updating code seems to be an overly exploitative phenomenon that high experimentation buffers, but without that exploitation, high experimentation stymies knowledge. This is further supported by the fact that high experimentation fares better with infrequent code updates than it does with no code updates. This suggests that having some degree of exploitation to balance it out yields better results. Hypothesis 2b is not supported since no experimentation performs roughly the same as low experimentation in disconnected structures. When the teams are disconnected, a small degree of experimentation seems to have little effect.

Without experimentation, organizational knowledge tends to peak when  $\beta = 0.1$  and then either level off or decline slightly. However, when there is no code as in the left panel of Figure 3, a small degree of experimentation offsets the decline that the non-experimental group has at  $\beta = 0.5$ . Not only does it prevent a decline, but it actually provides a small boost suggesting that the slight exploitation of higher inter-group connectivity is synergizing with experimentation to produce a better outcome.

Hypothesis 3a is only supported when there is high individual experimentation. At the other two levels of experimentation, a continuously updating code tends to have a negative effect on knowledge. The fact that knowledge is slightly higher at  $\beta = 0$  when the code updates infrequently (i.e.,  $\tau = 10$  in the right pane of Figure 3) than when there is no code suggests that offering some code learning to disconnected subgroups can have beneficial effects but that continuous updating is too exploitative. Hypothesis 3b is only supported when there is a small degree of experimentation since that yields the highest knowledge for tightly coupled subgroups with no code. However, when there is no experimentation, infrequent code updating performs



roughly the same. With high experimentation, non-existent code updates produce the lowest knowledge levels compared to other tightly coupled subgroups. Regarding intra-trial dynamics, sporadic updating yields a higher equilibrium than continuous updating. However, organizational knowledge shows a minor dip every time the code updates suggesting its overall undesirability in this model.

### **Exploratory Analyses**

When agents learned from each other and from an organizational code continuously, organizational knowledge suffered across conditions and interpersonal learning had little effect. However, when the code learning was removed (i.e., effectively,  $\tau = 0$ ), the dynamics from interpersonal learning became notable. The exploratory analyses ignore the organizational code altogether and focus on bottom-up learning. The most interesting patterns in this model emerged at the intersection of five key variables: (1) individual experimentation ( $\varepsilon$ ), (2) subgroup connectivity ( $\beta$ ), (3) interpersonal learning, (4) environmental turbulence ( $p_{env}$ ), and (5) turnover ( $p_{turn}$ ). Table 4 summarizes the parameter values tested in the exploratory analyses. The following sections highlight key findings. Changes to the tested parameter values in the exploratory analyses (e.g., number of agents and number of beliefs) were motivated by the desire to replicate March's (1991) model as closely as possible. This allows a reasonable comparison of his findings on the beneficial effects of turnover with individual experimentation in this bottom-up learning framework. Additionally, the maximum value of inter-team connection ( $\beta = 1$ ) produced knowledge levels similar to those at  $\beta = 0.2$  (or marginally lower as the 3 x 3 x 3 design revealed), so  $\beta = 0.2$  is the highest value included in results. Overall, the values of inter-team connection ( $\beta$ ) in the exploratory analyses were informed by sensitivity analyses such that more interesting shifts occurred when  $\beta$  was 0, 0.05, and 0.2. This also allowed for greater

parsimony in analyses.

### **Experimentation vs. No Experimentation**

Figure 4 shows the aggregate knowledge levels across the full range of interpersonal learning values ( $p_3$ ) and across three levels of  $\beta$  (i.e., 0, 0.05, and 0.2) without any environmental turbulence or turnover. Each plot shows trajectories with and without individual experimentation. When  $\beta = 0$ , the organization struggles to transmit useful knowledge beyond the strict team boundaries, and overall performance suffers dramatically. Adding inter-team connections tends to yield higher aggregate knowledge levels at  $\beta = 0.05$  and even higher at  $\beta = 0.2$ . Additionally, the slight inverted-U shape of each trajectory suggests that the optimal level of interpersonal learning falls somewhere between 0.2 and 0.8 in stable circumstances. Note that experimentation harms knowledge slightly at lower levels of  $p_3$  but improves it slightly at higher levels.

Figure 5, in comparison, shows the deleterious effect that changing reality ( $p_{env} = 0.02$ ) can have on aggregate knowledge levels. At every level of  $p_3$  and  $\beta$ , knowledge is lower with environmental turbulence than without it. With turbulence, experimentation raises the performance of completely isolated teams across every level of  $p_3$ . For other network structures, experimentation hinders performance slightly at very low levels of  $p_3$  but yields large improvements at higher levels and mitigates the dip seen at very high levels of  $p_3$  in Figure 4. Very high levels of interpersonal learning may be exploitative since random experimentation (i.e., exploration) yields higher knowledge levels compared to conditions without experimentation. Conversely, at very low levels of peer learning, information is not transferring well enough and this constitutes an overly exploratory strategy which preserves too much heterogeneity. Adding experimentation to those conditions only exacerbates the imbalance, albeit

slightly. Table 5 illustrates the interaction.

### **Experimentation vs. Turnover with Environmental Turbulence**

As mentioned previously, March (1991) showed the positive effects of turnover as variation in response to environmental turbulence. However, March's model had a formalized learning protocol in which individuals could only learn from the organizational code rather than each other. Figure 6 shows the effect of turnover ( $p_{turn} = 0.1$ ) in the face of environmental turbulence under this study's interpersonal learning protocol. Here, adding individual experimentation has almost no effect on knowledge at lower levels of interpersonal learning. However, at higher levels of interpersonal learning (i.e.,  $p_3 \geq 0.5$ ), adding individual experimentation improves aggregate knowledge by adding beneficial variance above and beyond turnover.

## DISCUSSION

Fang et al. (2010) argue in favor of the evolutionary perspective by showing that semi-isolated subgroups tend to strike the optimal balance between exploration and exploitation regardless of contingencies. However, one purpose of this study is to evaluate interventions that could potentially mitigate the deleterious effects of high and low subgroup coupling on organizational learning (and particularly the strategies that can combat unintended exploitation). This model examines the effects of individual experimentation and the frequency of organization-level communication of “best practices” on organizational learning through varying levels of network connectivity and environmental turbulence.

### **Tested Hypotheses**

Combining code learning and interpersonal learning had a large overall negative effect on organizational knowledge, and code learning muted the generally positive main effects of inter-team connectivity. That is, when the code updates at every time point, the positive effect of inter-team connection remains but the slope is lower than when the code updates sporadically or when the code is removed. It is possible that combining the two learning mechanisms was too exploitative such that learning from a code and from peers increased belief homogeneity and prevented proper adaptation to the environment. When employees learn from their most knowledgeable peers and from a code dictated by the organization’s most knowledgeable members, the two mechanisms are likely to reinforce the same opinions. When the code and the most knowledgeable peers are mostly correct in their beliefs, this might not be a problem. However, as these simulations indicate, strong socialization can also compound the exploitative effects of top-down learning.

When that exploitation was combined with a high degree of individual experimentation, however, those conditions yielded higher organizational knowledge than either no experimentation or low experimentation. In less exploitative conditions (i.e., when code learning was absent or infrequent), high levels of individual experimentation had a strongly harmful effect on organizational knowledge. This suggests that high individual experimentation can introduce too much random variation and is, therefore, too exploratory under those circumstances. Changing individual beliefs 10% of the time prevented valuable knowledge from accruing. This circumstance might correspond to an organization in which employees seek out a great deal of new and competing information. Maybe employees in this organization are too skeptical of the shared wisdom of their peers and even of their own past experiences. Switching beliefs erratically creates a failure trap in which organizational knowledge can never adequately vet beliefs that are so dynamic. However, as mentioned, this high degree of exploration proved useful in the otherwise exploitative condition in the middle pane of Figure 3. This result fits well with the notion that exploration and exploitation must be balanced to yield better knowledge outcomes.

The interaction between experimentation and network structure demonstrates the compensatory and even salutary effects that individual experimentation can have in highly connected organizations. Not only can individual experimentation counteract the exploitative tendency of highly connected networks, but it may also be able to synergize with it in some circumstances. For example, when there is no code and no experimentation, knowledge dips slightly going from loosely coupled to tightly coupled subgroups (in line with Fang et al., 2010). However, adding experimentation in this transition actually improves knowledge. This effect is not present in the right pane of Figure 3 when the code updates infrequently. It is possible that

the small degree of experimentation present was not sufficient to counteract the exploitation of code updating in that scenario. Future simulations should test a wider range of the experimentation parameter to determine if a higher degree of experimentation would replicate the beneficial effect of experimentation in the left pane of Figure 3.

### **Exploratory Simulations**

At very high levels of interpersonal learning, organizational knowledge tends to dip, presumably because the cognitive diversity disappears, and aggregate knowledge settles for a suboptimal equilibrium. Experimentation, however, tends to mitigate this exploitative and homogenizing effect of high interpersonal learning. Even with turnover injecting useful variance, individual experimentation can provide an advantage in cultures with high socialization rates, particularly in dynamic environments. This suggests that individually based trial-and-error may be a better strategy for adapting to environmental change than bringing in entirely new people. A small amount of healthy and enduring skepticism about one's own beliefs and the beliefs of peers can help organizations remain successful in tumultuous conditions.

Under the conditions tested in these exploratory simulations, the reward-to-risk ratio is high for individual experimentation. In a dynamic environment, experimentation allows organizational knowledge to keep pace and synergize with interpersonal learning. In a stable environment with relatively little interpersonal learning, experimentation can degrade performance slightly. When turnover is added to that circumstance, the deleterious effect of individual experimentation on organizational knowledge increases slightly. However, these negative effects are small, so the potential reward of experimentation in these circumstances seems to justify the minor risk across a variety of contexts. It should be noted, however, that the level of individual experimentation in these simulations is tempered. Here, each belief had a 1%

chance of changing at every time point. It is likely that significantly higher degrees of belief change would degrade organizational knowledge as one did in the first round of simulations. The degree of variance from experimentation that is desirable ultimately depends on the strength of opposing exploitative mechanisms.

Earlier discussion described some of the counterintuitive benefits of employee turnover (Staw, 1980). However, as an exploratory mechanism, individual experimentation is essentially cost-free for the organization. It delivers fresh perspectives but sidesteps the hassles of recruiting, selecting, onboarding, and training new employees. As a caveat, individual experimentation must be contained and orthogonal. In the evolutionary analogy previously discussed, individual experimentation parallels genetic mutation given its unpredictability. In nature, genetic mutations are harmful far more frequently than they are beneficial. The reason that mutation can be useful in this simulation is that the downsides of bad belief changes are highly constrained. When a belief changed in this simulation, it did so independent of any other event in the organization. Whole teams did not change beliefs simultaneously nor did the organization as a whole change tack randomly. When an individual agent changed a belief, the likelihood of that new belief spreading depended on the agent's influence (i.e., relative knowledge level among peers). Thus, the bottom-up learning mechanism constituted a check on new (and potentially bad) ideas. Additionally, the risk of changing any one belief (or conducting one mini-experiment) was constrained. To be an effective strategy in organizations, this type of bottom-up experimentation must carry negligible risks and may need to occur independently of pre-existing norms. While individual experimentation is cost-free, it may require some encouragement and cultural molding. Some level of psychological safety may be necessary for employees to feel comfortable trying new things at all and then to be able to share their findings

with coworkers and ultimately benefit the organization (Baer & Frese, 2003; Edmondson, 1999). These findings shed light on how organizations might try to balance exploration and exploitation in more modern, informal, bottom-up structures.

### **Limitations**

As mentioned previously, organizational learning has been studied with computational models because the mechanisms involved are often too complex to study empirically. Thus, this paper does no attempt to validate results through comparison to real world observations. Rather, Appendix C includes a discussion of how this model was validated against prior models in its lineage. Additionally, to construct a virtual simulation of organizational dynamics, all computational models require assumptions and simplifications (Harrison, Lin, Carroll, & Carley, 2007). This model makes assumptions similar to March's (1991) model and its extensions. What follows are highlights of the most significant ones.

Most notably, this framework abstracts the learning process in teams and organizations by drastically simplifying the mechanism of knowledge transmission and ignoring individual differences altogether. It assumes that agents can accurately assess the beliefs that superior agents hold. This is the primary way that agents change their own beliefs through interpersonal learning. However, research in the "hidden profile" paradigm has shown that people in groups tend to spend most time discussing shared knowledge (Stasser & Titus, 1985). While getting everyone on the same page is useful, failing to disclose unique and relevant information is both common and detrimental to team decision-making. It follows that real team members would not have access to every belief of every colleague to make fully informed decisions. Rather, the model assumes that agents use heuristics to infer how and why their esteemed peers perform better than themselves.



This model also assumes that knowledge emerges from individual to team level and from individual to organizational level in entirely compositional ways (Kozlowski & Klein, 2000). The beliefs of agents, teams, and the organizational code all share the same dimensional structure of reality. In *real* reality, however, humans categorize knowledge in different schemata, and previously held knowledge can influence how we encode new information (Anderson, 1977). Additionally, the model operationalizes belief change as immediate and total when actual change is more likely a gradual process of probability updating that can retain elements of what came before (Busemeyer & Townsend, 1993; Bohner & Dickel, 2011). Fiol and Lyles (1985) assert that “organizational learning is not simply the sum of each member’s learning,” so the operationalization of organizational knowledge as the average of individual knowledge is arguably imperfect. People hold beliefs configurations with far more nuance than simply positive, neutral, and negative. However, the structure and transmission of knowledge in March’s (1991) model are fundamental, and altering them would disqualify this model as an extension. This framework is incomplete but not entirely unrealistic, as compositional team cognition forms the basis of team mental models (Cannon-Bowers et al., 1993).

Lastly, there is no guarantee that more effective organizational learning will produce better performance at any level of an organization. Having a more accurate representation of reality should presumably fuel better decisions and competitive advantages as the Good Regulator Theorem suggests. However, while evidence for such a relationship certainly exists, it is not overwhelming (Jiménez- Jiménez & Sanz-Valle, 2011). This empirical dearth may be due to the difficulty of conducting research in this domain. In the theory of planned behavior (Ajzen, 1991), beliefs about the likely outcomes of behavior influence one’s attitude toward the behavior. This attitude along with subjective norms and perceived behavioral control relate to

behavioral intentions which finally relate to actual behavior. Thus, it may be important to separate conceptually organizational learning from the actions and decisions that organizational members make on the basis of their learned beliefs.

### **Future Research**

Because the reward-to-risk ratio of individual experimentation was favorable in the exploratory analyses, future research could focus on identifying the boundaries of that phenomenon. For example, at what level of experimentation does the exploratory benefit become maladaptive and impede the retention of valuable organizational knowledge? Do more turbulent environments require more individual experimentation in a linear fashion, or are there diminishing returns that might require other adaptive mechanisms to maintain sustainable knowledge?

Another promising avenue for future research is the role of compilational team cognition or transactive memory systems (Wegner, 1987) in inter-team and organizational learning. DeChurch and Mesmur-Magnus (2010) showed that while compositional team cognition positively impacts team performance, compilational cognition has a significantly stronger effect. Since compilational cognition is a distinctly team-level construct, it has more proximal influence on team-level performance. Rather than assuming that most or all people in a group need to share a piece of knowledge for it to emerge at a higher level, it would be valuable to develop theory through computation about how qualitatively different types of knowledge can remain distributed but accessible to everyone in a way that bolsters aggregate organizational knowledge.

Other computational models could use more complex experimentation parameters and at more than just the individual level. For example, one might assume that individuals explore based on a *dynamic allocation index* or *Gittins index* (Gittins, 1979). Like  $\epsilon$ -greedy, this emerged

as an attempt to deal with the multi-armed bandit problem. The decision of whether to exploit or explore is influenced by the interval of consideration. The Gittins index assumes an infinite future, as many organizations hope for, but it discounts future rewards given their lower certainty and lower desirability compared to immediate rewards. While this method is still limited (e.g., it uses linear discounting and does not account for switching costs), it could be an interesting and more nuanced way to operationalize team and organizational experimentation going forward.

Moore's law suggests that technology grows exponentially rather than linearly (Moore, 1965). The current model assumes that each dimension of reality has a fixed probability of switching its value at each time point. However, because Moore's prediction seems to have held true even longer than he anticipated, future models might experiment with rates of environmental change that increase over time (e.g., doubling every 18 months or every 2 years as theories suggest). Similarly, Harrison & Carroll (1991) included organizational growth rate (i.e., increasing total number of employees) and selectivity (i.e., recruiting people with varying degrees of similarity to the organizational code). They also tested different configurations of the core parameters to mimic different organizational types (e.g., an entrepreneurial model with high growth rates or a Japanese business model with an intensive socialization and very low turnover).

Future models could employ even more realistic organizational structures with multi-tiered hierarchies (e.g., building from Bray & Prietula, 2007) and specialized roles. They might also vary the degree to which individuals can shape the organizational code. Given the importance of specialization and division of labor in organizations, it could be useful to specify distinct but overlapping knowledge sets for different roles or individuals. Roles can dictate where people look for information (Simon, 1991), and not all information is relevant to every person.

Varying roles might synergize with the aforementioned exploration of transactive memory systems as well. Additionally, future models might examine multiple goals by supplementing the current selecting mechanism (i.e., knowledge or concordance with reality) with a desire to act in alignment with the organizational code. When the dimensions of reality and the organizational code are not perfectly aligned, employees may experience a tension between adapting to the external environment and adapting to the internal environment (i.e., the organizational culture).

### **Practical Implications**

To demonstrate the utility of this model and the consideration of the exploration/exploitation tradeoff, consider *Valve*, a video game design company based in Bellevue, Washington that purports to have an entirely flat organizational structure. They publish their employee handbook (Valve, 2012) which describes the nuances of their non-traditional non-hierarchy. While it is possible that their espoused and enacted values differ (Zohar & Hofmann, 2012), their philosophy revolves around the autonomy to get involved in multiple and varied team-based projects throughout the company as well as learning from a great deal of informal interaction rather than top-down training. In the context of this proposed model, Valve shows relatively tight subgroup coupling or high organizational connectivity. This contemporary structure might have exploratory intentions, but this paper has discussed the covertly homogenizing effect that it can have.

It appears, however, that Valve has some strategies that might counterbalance its unintentionally exploitative structure. The handbook explicitly states that their structure makes them bad at disseminating information internally and that mentoring newcomers is not a strength. This suggests that their organizational code requires a great deal of tacit learning. Their handbook expresses a deep respect for individuality and encourages healthy disagreement among

colleagues. It recommends scrutinizing the founder/president's ideas as heavily and honestly as any other coworker's. These notions combat homogeneity by guarding against socialization that is too swift. The handbook also asserts that, "Screwing up is a great way to find out that your assumptions were wrong or that your model of the world was a little bit off." This suggests that individual experimentation and learning through failure are not only tolerated but also encouraged at Valve. This exploratory maxim might infuse needed variance into a system that might otherwise be pointed toward exploitation. The "managerial" strategies at the Valve Corporation illustrate how a modern organization might balance exploration and exploitation.

## CONCLUSION

Because organizational learning is such a complex process to capture empirically, computational modeling has proven to be a useful analytical and exploratory tool. March's (1991) model sparked conversation about exploration and exploitation in organizational science and strategy that continues today. While this model is computational and abstract, it nonetheless clarifies some specific circumstances and concrete strategies that can help individuals, teams, and organizations in non-hierarchical networks balance exploration and exploitation and set themselves up for long-term learning.

## APPENDICES

## APPENDIX A

### Tables and Figures



Figure 1. Notional Representation of the Baseline Network Structure ( $\beta = 0$ )





Figure 2. Notional Representation of How Network Ties are Replaced



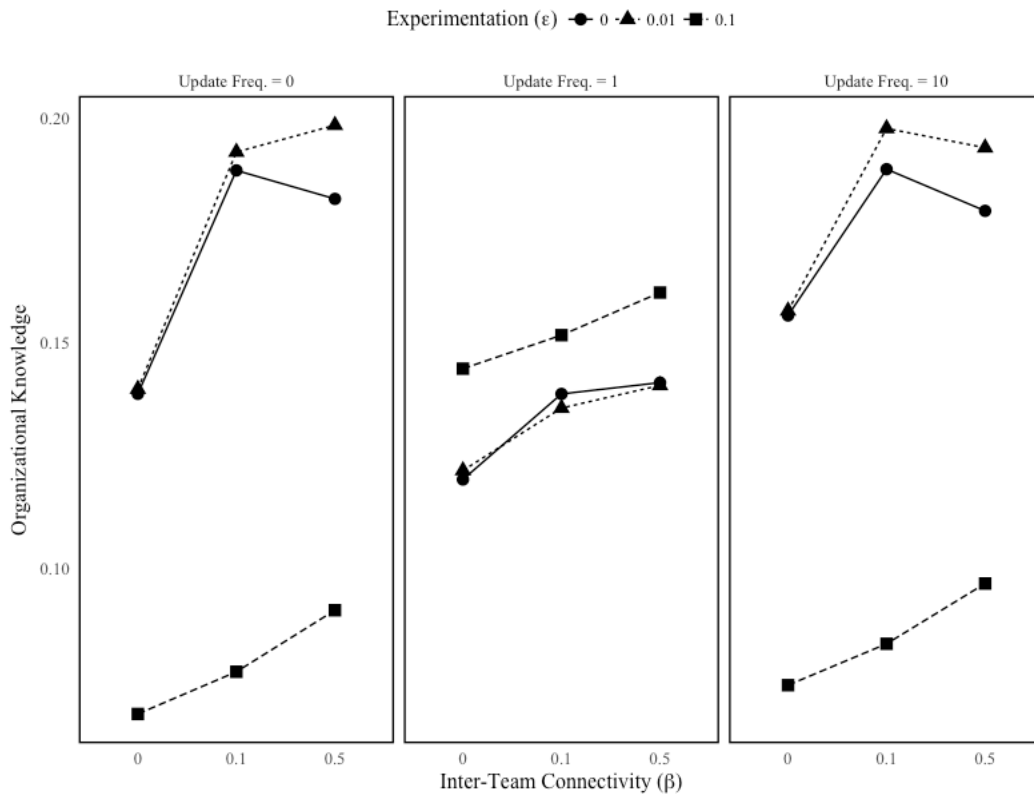
*Table 1. Notional Representation of Initial Conditions*

|             | Belief 1 | Belief 2 | Belief 3 | Knowledge |
|-------------|----------|----------|----------|-----------|
| Environment | -1       | 1        | 1        | NA        |
| Agent 1     | -1       | 1        | 0        | 2         |
| Agent 2     | 1        | 0        | 1        | 0         |
| Agent 3     | 0        | -1       | -1       | -2        |
| Code        | 0        | 0        | 0        | 0         |

*Table 2. Model Parameters and their Simulated Values in 3x3x3 Design*

| Parameters    | Description   | Range of parameter values analyzed |
|---------------|---|------------------------------------|
| $n$           | Number of individuals in the organization                 | 140                                |
| $m$           | Number of belief dimensions                               | 100                                |
| $z$           | Size of subgroup  | 7                                  |
| $q$           | Number of dimensions in the code that are tacit           | 50                                 |
| $\beta$       | Probability of adding social ties to the network          | 0, 0.1, 0.5                        |
| $p_1$         | Probability of learning from the organizational code      | 0.1                                |
| $p_2$         | Probability of the code learning from PMEs                | 0.9                                |
| $p_3$         | Probability of learning from peers                        | 0.3                                |
| $p_{turn}$    | Probability of an agent being replaced with a naïve agent | 0.1                                |
| $p_{envir}$   | Probability of reality dimensions changing                | 0.02                               |
| $\varepsilon$ | Probability of an agent changing a belief                 | 0, 0.01, 0.1                       |
| $\tau$        | Number of time periods between code updates               | 0 (i.e., no code), 1, 10           |
| $T$           | Number of time periods per trial                          | 200                                |
| $N$           | Number of runs per condition                              | 200                                |

Figure 3. Organizational Knowledge in the 3x3x3 Design



Note. Inter-Team Connectivity values represented on the horizontal axes are spaced evenly, not by magnitude.

Table 3. Effect Sizes Between All Conditions in the 3x3x3 Design

| $\varepsilon$ | $\tau$ | $\beta$ |    |    |    |     |     |     |    |    |    |  |
|---------------|--------|---------|----|----|----|-----|-----|-----|----|----|----|--|
| $\varepsilon$ |        |         | 0  | 0  | 0  | .01 | .01 | .01 | .1 | .1 | .1 |  |
| $\tau$        |        |         | 0  | 0  | 0  | 0   | 0   | 0   | 0  | 0  | 0  |  |
| $\beta$       |        |         | 0  | .1 | .5 | 0   | .1  | .5  | 0  | .1 | .5 |  |
|               | 0      | 0       | 0  |    |    |     |     |     |    |    |    |  |
|               | 0      | 0       | .1 |    |    |     |     |     |    |    |    |  |
|               | 0      | 0       | .5 |    |    |     |     |     |    |    |    |  |
|               | .01    | 0       | 0  |    |    |     |     |     |    |    |    |  |
|               | .01    | 0       | .1 |    |    |     |     |     |    |    |    |  |
|               | .01    | 0       | .5 |    |    |     |     |     |    |    |    |  |
|               | .1     | 0       | 0  |    |    |     |     |     |    |    |    |  |
|               | .1     | 0       | .1 |    |    |     |     |     |    |    |    |  |
|               | .1     | 0       | .5 |    |    |     |     |     |    |    |    |  |
|               | 0      | 1       | 0  |    |    |     |     |     |    |    |    |  |
|               | 0      | 1       | .1 |    |    |     |     |     |    |    |    |  |
|               | 0      | 1       | .5 |    |    |     |     |     |    |    |    |  |
|               | .01    | 1       | 0  |    |    |     |     |     |    |    |    |  |
|               | .01    | 1       | .1 |    |    |     |     |     |    |    |    |  |
|               | .01    | 1       | .5 |    |    |     |     |     |    |    |    |  |
|               | .1     | 1       | 0  |    |    |     |     |     |    |    |    |  |
|               | .1     | 1       | .1 |    |    |     |     |     |    |    |    |  |
|               | .1     | 1       | .5 |    |    |     |     |     |    |    |    |  |
|               | 0      | 10      | 0  |    |    |     |     |     |    |    |    |  |
|               | 0      | 10      | .1 |    |    |     |     |     |    |    |    |  |
|               | 0      | 10      | .5 |    |    |     |     |     |    |    |    |  |
|               | .01    | 10      | 0  |    |    |     |     |     |    |    |    |  |
|               | .01    | 10      | .1 |    |    |     |     |     |    |    |    |  |
|               | .01    | 10      | .5 |    |    |     |     |     |    |    |    |  |
|               | .1     | 10      | 0  |    |    |     |     |     |    |    |    |  |
|               | .1     | 10      | .1 |    |    |     |     |     |    |    |    |  |
|               | .1     | 10      | .5 |    |    |     |     |     |    |    |    |  |

Note. Positive values indicate that the condition described on the horizontal axis had higher aggregate knowledge across trials.

Table 3 (continued)

| $\varepsilon$ | $\tau$ | $\beta$ | 0     | 0     | 0     | 0.01  | 0.01  | 0.01  | 0.1    | 0.1    | 0.1    |
|---------------|--------|---------|-------|-------|-------|-------|-------|-------|--------|--------|--------|
| $\varepsilon$ | $\tau$ | $\beta$ | 1     | 1     | 1     | 1     | 1     | 1     | 1      | 1      | 1      |
| $\varepsilon$ | $\tau$ | $\beta$ | 0     | 0.1   | 0.5   | 0     | 0.1   | 0.5   | 0      | 0.1    | 0.5    |
| 0             | 0      | 0       |       |       |       |       |       |       |        |        |        |
| 0             | 0      | .1      |       |       |       |       |       |       |        |        |        |
| 0             | 0      | .5      |       |       |       |       |       |       |        |        |        |
| .01           | 0      | 0       |       |       |       |       |       |       |        |        |        |
| .01           | 0      | .1      |       |       |       |       |       |       |        |        |        |
| .01           | 0      | .5      |       |       |       |       |       |       |        |        |        |
| .1            | 0      | 0       |       |       |       |       |       |       |        |        |        |
| .1            | 0      | .1      |       |       |       |       |       |       |        |        |        |
| .1            | 0      | .5      |       |       |       |       |       |       |        |        |        |
| 0             | 1      | 0       | -     |       |       |       |       |       |        |        |        |
| 0             | 1      | .1      | 0.9   | -     |       |       |       |       |        |        |        |
| 0             | 1      | .5      | 1.02  | 0.11  | -     |       |       |       |        |        |        |
| .01           | 1      | 0       | 0.11  | -0.84 | -0.96 | -     |       |       |        |        |        |
| .01           | 1      | .1      | 0.76  | -0.14 | -0.25 | 0.69  | -     |       |        |        |        |
| .01           | 1      | .5      | 1.05  | 0.09  | -0.03 | 1     | 0.24  | -     |        |        |        |
| .1            | 1      | 0       | 1.71  | 0.35  | 0.19  | 1.73  | 0.54  | 0.25  | -      |        |        |
| .1            | 1      | .1      | 2.23  | 0.81  | 0.65  | 2.3   | 1.01  | 0.76  | 1.42   | -      |        |
| .1            | 1      | .5      | 2.86  | 1.38  | 1.22  | 2.99  | 1.58  | 1.4   | 3      | 1.7    | -      |
| 0             | 10     | 0       | 2.42  | 1.04  | 0.88  | 2.5   | 1.23  | 1.02  | 1.72   | 0.64   | -0.72  |
| 0             | 10     | .1      | 4.28  | 2.81  | 2.65  | 4.48  | 3     | 2.94  | 4.93   | 4.13   | 2.99   |
| 0             | 10     | .5      | 3.55  | 2.21  | 2.06  | 3.68  | 2.39  | 2.28  | 3.45   | 2.73   | 1.76   |
| .01           | 10     | 0       | 2.53  | 1.12  | 0.96  | 2.63  | 1.31  | 1.1   | 2.04   | 0.87   | -0.62  |
| .01           | 10     | .1      | 4.94  | 3.38  | 3.21  | 5.2   | 3.57  | 3.57  | 6.36   | 5.51   | 4.25   |
| .01           | 10     | .5      | 4.49  | 3.04  | 2.87  | 4.69  | 3.22  | 3.18  | 5.16   | 4.4    | 3.32   |
| .1            | 10     | 0       | -3.26 | -4.08 | -4.19 | -3.76 | -3.9  | -4.67 | -16.66 | -18.96 | -18.99 |
| .1            | 10     | .1      | -2.6  | -3.49 | -3.61 | -3.03 | -3.31 | -4.02 | -14.21 | -16.38 | -16.72 |
| .1            | 10     | .5      | -1.64 | -2.64 | -2.77 | -1.97 | -2.46 | -3.07 | -10.63 | -12.6  | -13.35 |

Note. Positive values indicate that the condition described on the horizontal axis had higher aggregate knowledge across trials.

Table 3 (continued)

| $\varepsilon$ | $\tau$ | $\beta$ | 0  | 0      | 0      | 0.01   | 0.01   | 0.01   | 0.1    | 0.1  | 0.1  |   |
|---------------|--------|---------|----|--------|--------|--------|--------|--------|--------|------|------|---|
| $\varepsilon$ |        |         | 0  | 0      | 0      | 0.01   | 0.01   | 0.01   | 0.1    | 0.1  | 0.1  |   |
| $\tau$        |        |         | 10 | 10     | 10     | 10     | 10     | 10     | 10     | 10   | 10   |   |
| $\beta$       |        |         | 0  | 0.1    | 0.5    | 0      | 0.1    | 0.5    | 0      | 0.1  | 0.5  |   |
|               | 0      | 0       | 0  |        |        |        |        |        |        |      |      |   |
|               | 0      | 0       | .1 |        |        |        |        |        |        |      |      |   |
|               | 0      | 0       | .5 |        |        |        |        |        |        |      |      |   |
|               | .01    | 0       | 0  |        |        |        |        |        |        |      |      |   |
|               | .01    | 0       | .1 |        |        |        |        |        |        |      |      |   |
|               | .01    | 0       | .5 |        |        |        |        |        |        |      |      |   |
|               | .1     | 0       | 0  |        |        |        |        |        |        |      |      |   |
|               | .1     | 0       | .1 |        |        |        |        |        |        |      |      |   |
|               | .1     | 0       | .5 |        |        |        |        |        |        |      |      |   |
|               | 0      | 1       | 0  |        |        |        |        |        |        |      |      |   |
|               | 0      | 1       | .1 |        |        |        |        |        |        |      |      |   |
|               | 0      | 1       | .5 |        |        |        |        |        |        |      |      |   |
|               | .01    | 1       | 0  |        |        |        |        |        |        |      |      |   |
|               | .01    | 1       | .1 |        |        |        |        |        |        |      |      |   |
|               | .01    | 1       | .5 |        |        |        |        |        |        |      |      |   |
|               | .1     | 1       | 0  |        |        |        |        |        |        |      |      |   |
|               | .1     | 1       | .1 |        |        |        |        |        |        |      |      |   |
|               | .1     | 1       | .5 |        |        |        |        |        |        |      |      |   |
|               | 0      | 10      | 0  | -      |        |        |        |        |        |      |      |   |
|               | 0      | 10      | .1 | 3.27   | -      |        |        |        |        |      |      |   |
|               | 0      | 10      | .5 | 2.11   | -0.74  | -      |        |        |        |      |      |   |
|               | .01    | 10      | 0  | 0.14   | -3.28  | -2.08  | -      |        |        |      |      |   |
|               | .01    | 10      | .1 | 4.41   | 0.82   | 1.52   | 4.49   | -      |        |      |      |   |
|               | .01    | 10      | .5 | 3.57   | 0.4    | 1.09   | 3.59   | -0.37  | -      |      |      |   |
|               | .1     | 10      | 0  | -13.64 | -13.71 | -10.97 | -15.45 | -16.02 | -13.37 | -    |      |   |
|               | .1     | 10      | .1 | -12    | -12.55 | -9.97  | -13.58 | -14.75 | -12.29 | 3.3  | -    |   |
|               | .1     | 10      | .5 | -9.58  | -10.83 | -8.51  | -10.83 | -12.85 | -10.68 | 7.36 | 4.21 | - |

Note. Positive values indicate that the condition described on the horizontal axis had higher aggregate knowledge across trials.

Table 4. Parameters in Exploratory Analyses and their Simulated Values

| Parameters    | Description   | Range of parameter values analyzed |
|---------------|---|------------------------------------|
| $n$           | Number of individuals in the organization                 | 50                                 |
| $m$           | Number of belief dimensions                               | 30                                 |
| $z$           | Size of subgroup  | 5                                  |
| $\beta$       | Probability of adding social ties to the network          | 0, 0.05, 0.2                       |
| $p_3$         | Probability of learning from peers                        | 0 through 1 in increments of 0.1   |
| $p_{turn}$    | Probability of an agent being replaced with a naïve agent | 0, 0.1                             |
| $p_{env}$     | Probability of reality dimensions changing                | 0, 0.02                            |
| $\varepsilon$ | Probability of an agent changing a belief                 | 0, 0.01                            |
| $T$           | Number of time periods per trial                          | 100                                |
| $N$           | Number of runs per condition                              | 80                                 |

Figure 4. Organizational Knowledge with No Turbulence and No Turnover




Figure 5. Organizational Knowledge with Turbulence ( $p_{env} = 0.02$ ) and No Turnover  
 Experimentation ( $\epsilon$ )  $\bullet$  0  $\blacktriangle$  0.01




Table 5. Ranges of  $p_3$  in which Experimentation is Better than No Experimentation

| $p_{turn}$ | $p_{env}$ | $\beta$ | $p_3$ Range    |
|------------|-----------|---------|----------------|
| 0          | 0         | 0       | $p_3 > 0$      |
|            |           | 0.05    | $p_3 \geq 0.3$ |
|            |           | 0.2     | $p_3 \geq 0.8$ |
| 0          | 0.02      | 0       | $p_3 \geq 0.1$ |
|            |           | 0.05    | $p_3 \geq 0.2$ |
|            |           | 0.2     | $p_3 \geq 0.2$ |
| 0.1        | 0         | 0       | $p_3 \geq 0.6$ |
|            |           | 0.05    | $p_3 \geq 0.9$ |
|            |           | 0.2     | $p_3 = 1$      |
| 0.1        | 0.02      | 0       | $p_3 \geq 0.5$ |
|            |           | 0.05    | $p_3 \geq 0.6$ |
|            |           | 0.2     | $p_3 \geq 0.7$ |

Note. Ranges indicate the levels of  $p_3$  at which experimentation (i.e.,  $\epsilon = 0.01$ ) yields significantly ( $p < 0.05$ ) higher aggregate knowledge than with no experimentation (i.e.,  $\epsilon = 0$ ).



Figure 6. Organizational Knowledge with  
 Both Turbulence ( $p_{env} = 0.02$ ) and Turnover ( $p_{turn} = 0.1$ )  
 Experimentation ( $\epsilon$ )  $\bullet$  0  $\blacktriangle$  0.01






Figure 7. Replication of March (1991)



Note. Parameter values are  $m = 30$ ,  $n = 50$ ,  $p_1 = 0.5$ ,  $p_2 = 0.5$ ,  $p_{env} = 0.02$ ,  $N = 80$ .

Figure 8. Conceptual Replication of Fang et al. (2010)



Note. This represents average organizational knowledge while Fang et al. plotted average equilibria. The results are still roughly similar, but when  $\beta = 0$ , knowledge is lower in this model than it was in Fang et al. This is likely due to subgroups being completely isolated in the current model whereas the baseline in Fang et al. had some intergroup ties.

## APPENDIX B

### R Simulation Code

```

#####
# Define Simulation Function
#####

full_sim <- function(n, m, z, q, beta, p1, p2, p3, p.turn, p.env, epsilon, tau, T, N){

  # Initialize vectors & matrices
  varnames <- c('beta', 'q', 'p1', 'p2', 'p3', 'p.turn', 'p.env', 'epsilon', 'tau')
  num.conditions <- length(beta)*length(q)*length(p1)*length(p2)*length(p3)*
    length(p.turn)*length(p.env)*length(epsilon)*length(tau)
  avg.know.mat <- matrix(as.numeric(NA), nrow = N, ncol = T)
  avg.know.trial <- rep(as.numeric(NA), N)
  aggregate.know <- rep(as.numeric(NA), num.conditions)

  # Specify conditions
  if(num.conditions == 1) {
    combinations <- t(as.matrix(sapply(varnames, function(x) get(x) )))
  } else {
    combinations <- expand.grid(sapply(varnames, function(x) get(x) ))
  }

  # Create output matrix
  p.matrix <-
    cbind(combinations,
          matrix(as.numeric(NA), nrow = num.conditions, ncol = 1 + N,
                dimnames = list(NULL, c('aggregate.know', paste0("N", 1:N)))))

  #####
  # Define Sub-Functions
  #####

  # Environmental turbulence
  poss.env <- c(-1L, 1L)
  turbulence <- function(x){
    return(poss.env[poss.env != x])
  }

  # Code Learning: Determine the probability of the code learning
  get.p <- function(x){
    return(1 - (1 - p.matrix[cond.num, 'p2'])^k[x])
  }

  # Individual experimentation
  poss.beliefs <- c(-1L, 1L)
  experiment <- function(x){
    return(sample(poss.beliefs[poss.beliefs != x], 1))
  }
}

```

```

}

# Peer Learning: Determine from whom and what each agent can learn
what.to.learn <- function(x){
  peers.to.learn.from = which(net.know[x, ])
  dims.to.learn = peer.learn.list[[which(will.learn.vec == x)]]
  sup.belief.matrix = d[peers.to.learn.from, dims.to.learn, drop = F]
  sup.belief.vector = colSums(sup.belief.matrix)
  sup.belief.vector[sup.belief.vector > 0] = 1
  sup.belief.vector[sup.belief.vector < 0] = -1
  return(sup.belief.vector)
}

# Peer Learning: Helps make peer.learn.list
make.peer.list <- function(x){
  return(which(peer.learn[x, ]))
}

# Faster mean function
avg <- function(x){
  sum(x, na.rm = T) / length(x)
}

#####
# Loops
#####

for(cond.num in 1:nrow(p.matrix)){ # Begin Condition Loop -----
  for(trial.num in 1:N){ # Begin Trial Loop -----

    #####
    # Create Network Matrix
    #####

    # Baseline matrix
    network <- matrix(0L, nrow = n , ncol = n)
    teams <- matrix(1:n, nrow = z, ncol = n/z)
    for(i in 1:(n/z)){
      network[teams[, i], teams[, i]] <- 1L
    }
    diag(network) <- 0

    # Remove and add ties to baseline according to beta
    net.switch <- matrix(NA, nrow = n, ncol = n)
    for(i in 1:n){
      rnums <- runif(sum(network[i, ]), 0, 1) < p.matrix[cond.num, 'beta']

```

```

if(sum(rnums) == 0){next}
net.switch[i, ] <- network[i, ] == 1
net.switch[i, which(net.switch[i, ])] <- rnums
local.team <- teams[, as.numeric(which(teams == i, arr.ind = T)[, 2])]
avail.cons <- (1:n)[as.logical((network == 0)[i, ]*((1:n)!=local.team))]
newcon <- sample(avail.cons, sum(rnums))
network[i, newcon] <- 1
network[newcon, i] <- 1
column <- which(net.switch[i, ])
network[i, column] <- 0
network[column, i] <- 0
}

# Set diagonal to NA and make network matrix logical
diag(network) <- NA
network <- network == 1
network[network == FALSE] <- NA

# Working matrix for peer learning
net.know <- matrix(0, nrow = n , ncol = n)

#####
# Initial conditions for each trial
#####

d <- matrix(sample(c(-1L, 0L, 1L), n*m, replace = TRUE), nrow = n, ncol = m)
explicit_dims <- m - p.matrix[cond.num, 'q']
code <- rep(0, explicit_dims)
env <- sample(c(-1L, 1L), m, replace = TRUE)
know <- as.vector(d%*%env)
avg.know <- rep(as.numeric(NA), T)
code.know <- as.numeric(c(0, rep(NA, T - 1)))
PME <- know > code.know[1]
PME[!PME] <- NA
no.code <- p.matrix[cond.num, 'tau'] == 0
PMEbeliefs <- PME*d[, 1:explicit_dims]

for(time.pt in 1:T){ # Begin Time Loop -----

# Code knowledge level (at each T)
if(time.pt > 1){
  code.know[time.pt] <- sum(code * env[1:explicit_dims])
}

# Recalculate individual knowledge levels
know <- as.vector(d%*%env)

```

```

# Who's in the policy making elite?
PME <- know > code.know[time.pt]
PME[!PME] <- NA

# Average individual knowledge (at each T)
avg.know[time.pt] <- avg(know)

# Learning BY the code
PMEbeliefs <- PME*d[, 1:explicit_dims]
if(no.code == FALSE & time.pt%%p.matrix[cond.num, 'tau'] == 0){
  which.match.code <- PMEbeliefs == matrix(rep(code,each = n), nrow=n)
  k <- apply(which.match.code, 2, function(x){
    do.not.differ = sum(x, na.rm = T)
    differ = sum(!x, na.rm = T)
    output = differ - do.not.differ
    return(ifelse(output > 0, output, 0))
  })
  what.code.lerns <- sapply(1:length(k), function(x){
    if(k[x] > 0){
      position = sum(PMEbeliefs[which(PMEbeliefs[, x] != code[x]), x])
      return(ifelse(position > 0, 1, ifelse(position < 0, -1, 0)))
    } else {
      return(NA)
    }
  })
  p.by.code <- sapply(1:length(code), get.p)
  by.code <- runif(explicit_dims, 0, 1)
  by.code <- code != what.code.lerns & by.code < p.by.code
  code[by.code] <- what.code.lerns[by.code]
}

# Individual experimentation
exp <- matrix(runif(n*m, 0, 1), nrow = n, ncol = m) < p.matrix[cond.num, 'epsilon']
if (sum(exp) > 0) {
  toexp <- d[exp]
  d[exp] <- ifelse(toexp == 1, -1,
    ifelse(toexp == -1, 1, sample(poss.beliefs, 1)))
}

# Interpersonal learning
peer.learn <- matrix(runif(n*m, 0, 1), nrow = n, ncol = m) < p.matrix[cond.num, 'p3']
selected.to.learn <- rowSums(peer.learn) > 0
net.know <- t(network * know) > know
net.know[net.know == FALSE] <- NA
has.sups <- rowSums(net.know, na.rm = T) > 0

```



```

will.learn.vec <- which(selected.to.learn & has.sups)
peer.learn.list <- lapply(will.learn.vec, function(x){which(peer.learn[x, ])}))
peer.learn[-will.learn.vec, ] <- FALSE
replacements <- unlist(lapply(will.learn.vec, what.to.learn))
if (!is.null(replacements)){
  peer.learn <- t(peer.learn)
  d <- t(d)
  d[peer.learn] <- replacements
  peer.learn <- t(peer.learn)
  d <- t(d)
}

# Learning FROM the code
if(no.code == FALSE & time.pt%%p.matrix[cond.num, 'tau'] == 0){
  fc1 <- t(sapply(1:nrow(d),
    function(x){return(d[x, 1:explicit_dims] != code & code != 0)}))
  fc2 <- matrix(as.numeric(NA), nrow = n, ncol = explicit_dims)
  fc2[fc1] <- runif(sum(fc1), 0, 1)
  fc2 <- fc2 < p.matrix[cond.num, 'p1']
  fc2[is.na(fc2)] <- FALSE
  if (any(fc2)){
    d[fc2] <- unlist(lapply(1:explicit_dims, function(x){rep(code[x], sum(fc2[, x]))}))
  }

# Turnover
turn <- c(runif(n, 0, 1)) < p.matrix[cond.num, 'p.turn']
d[turn,] <- sample(c(-1, 0, 1), m*sum(turn), replace = TRUE)

# Environmental turbulence
env.change <- c(runif(m, 0, 1)) < p.matrix[cond.num, 'p.env']
env[env.change] <- unlist(sapply(env[env.change], turbulence))

} # End Time Loop -----

avg.know.mat[trial.num, ] <- avg.know
avg.know.trial[trial.num] <- mean(avg.know) / m

} # End Trial Loop -----

avg.know.by.time <- colMeans(avg.know.mat) / m
p.matrix[cond.num, c('aggregate.know', paste0("N", 1:N))] <-
  c(avg(avg.know.by.time), avg.know.trial)

} # End Condition Loop -----
return(p.matrix)
}

```

## APPENDIX C

### Replication Discussion

Not only does this model mimic certain aspects of March (1991) and Fang et al. (2010), but it also compares some key findings. For this reason, it is important to demonstrate that the current incarnation of the model has faithfully replicated the important aspects of prior models. What follows is a description of those replication efforts.

The ways in which this model draws from March (1991) have already been discussed. However, one of March's findings is particularly relevant to this model. He found that turnover allowed organizational knowledge (operationalized as the concordance of the organizational code with reality) to maintain a relatively high equilibrium in the face of environmental turbulence. Without that turnover, organizational knowledge would rise quickly but then fall dramatically as the organization failed to adapt to changing circumstances. Using the same parameter values in the current model yielded a satisfactory replication as shown in Figure 7.

Hypothesis 1 in this study seeks a conceptual replication of a finding from Fang et al. (2010). Namely that isolated subgroups fare poorly, loosely coupled subgroups perform best, and further inter-group connectivity beyond a low level degrades knowledge very slightly. Because the network structure in this model is not exactly the same as in the original, the standards for replication are less stringent. Results from this model (Figure 8) show roughly the same pattern as Fang et al.: a steep ascent from isolated subgroups to the peak and then a gradual and subtle descent as groups become more interconnected.

## REFERENCES

## REFERENCES

- Anderson, R. C. (1977). The notion of schemata and the educational enterprise: General discussion of the conference. In R. C. Anderson, R. J. Spiro, and W. E. Montague (Eds.) *Schooling and the Acquisition of Knowledge*, (pp. 415-431). Hillsdale, NJ: Erlbaum.
- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2), 179-211.
- Argote, L., Insko, C. A., Yovetich, N., & Romero, A. A. (1995). Group learning curves: The effects of turnover and task complexity on group performance. *Journal of Applied Social Psychology*, 25(6), 512-529.
- Argyris, C. & Schön, D. A. (1978). *Organizational learning: A theory of action perspective*. Addison-Wesley Publishing Co.
- Ashby, W. R. (1956). *An introduction to cybernetics*. London: Chapman & Hall LTD.
- Baer, M. & Frese, M. (2003). Innovation is not enough: Climates for initiative and psychological safety, process innovations, and firm performance. *Journal of Organizational Behavior*, 24(1), 45-68.
- Bandura, A., & Walters, R. H. (1963). *Social learning and personality development*. New York: Holt, Rinehart and Winston.
- Bantel, K. A. & Jackson S. E. (1989). Top management and innovations in banking: Does the composition of the top team make a difference? *Strategic Management Journal*, 10(1), 107-124.
- Baron, J. (2000). *Thinking and deciding*. Cambridge University Press.
- Bell, B. S., & Kozlowski, S. W. (2008). Active learning: Effects of core training design elements on self-regulatory processes, learning, and adaptability. *Journal of Applied Psychology*, 93(2), 296.
- Bell, S. T., Villado, A. J., Lukasik, M. A., Belau, L., & Briggs, A. L. (2010). Getting specific about demographic diversity variable and team performance relationships: A meta-analysis. *Journal of Management*, 37(3), 709-743.
- Benner, M. J., & Tushman, M. L. (2003). Exploitation, exploration, and process management: The productivity dilemma revisited. *Academy of Management Review*, 28(2), 238-256.
- Bennis, W. G. (1967). The coming death of bureaucracy. *Journal of Occupational and Environmental Medicine*, 9(7), 380.

- Bohner, G. & Dickel, N. (2011). Attitudes and attitude change. *Annual Review of Psychology*, 62, 391-417.
- Bray, D. A., & Prietula, M. J. (2007). Extending March's exploration and exploitation: Managing knowledge in turbulent environments. In *Twenty Eighth International Conference on Information Systems* (pp. 1-17).
- Burns, T. E., & Stalker, G. M. (1961). *The management of innovation*. London: Tavistock Publications.
- Busemeyer, J. R. & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3), 432-459.
- Buss, D. (2012). *Evolutionary psychology: The new science of the mind*. Psychology Press.
- Cannon-Bowers, J. A., Salas, E., & Converse, S. A. (1993). Shared mental models in expert team decision making. In J. N. J. Castellan (Ed.), *Current issues in individual and group decision making*. Hillsdale, NJ: Erlbaum.
- Cleeremans, A., & Dienes, Z. (2008). Computational models of implicit learning. In R. Sun (Ed.) *The Cambridge Handbook of Computational Psychology*, 396-421.
- Clune, J., Mouret, J. B., & Lipson, H. (2013). The evolutionary origins of modularity. *Proceedings of the Royal Society B: Biological Sciences*, 280(1755).
- Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), 128-152.
- Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89-97.
- Črepinšek, M., Liu, S. H., & Mernik, M. (2013). Exploration and exploitation in evolutionary algorithms: A survey. *ACM Computing Surveys (CSUR)*, 45(3), 35.
- Davis, G. F. (2016). *The vanishing American corporation: Navigating the hazards of a new economy*. Berrett-Koehler Publishers.
- Davis, J. H., Kerr, N. L., Atkin, R. S., Holt, R., & Meek, D. (1975). The decision processes of 6- and 12-person mock juries assigned unanimous and two thirds majority rules. *Journal of Personality and Social Psychology*, 32, 1-14.
- Dawkins, R. (2006). *The selfish gene*. Oxford University Press.
- De Dreu, C. K. W. & West, M. A. (2001). Minority dissent and team innovation: The importance

- of participation in decision making. *Journal of Applied Psychology*, 86(6), 1191-1201.
- DeChurch, L. A., Mesmer-Magnus, J. R. (2010). The cognitive underpinnings of effective teamwork: A meta-analysis. *Journal of Applied Psychology*, 95(1), 32-53.
- Dodgson, M. (1993). Organizational learning: A review of some literatures. *Organization Studies*, 14(3), 375-394.
- Edmondson, A. (1999). Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*, 44(2), 350-383.
- Eiben, A. E., & Schippers, C. A. (1998). On evolutionary exploration and exploitation. *Fundamenta Informaticae*, 35(1-4), 35-50.
- Fang, C., Lee, J., & Schilling, M. A. (2010). Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organization Science*, 21(3), 625-642.
- Fiol, C. M., & Lyles, M. A. (1985). Organizational learning. *Academy of Management Review*, 10(4), 803-813.
- Gittins, J. C. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)*, 148-177.
- Grand, J. A., Braun, M. T., Kuljanin, G., Kozlowski, S. W., & Chao, G. T. (2016). The dynamics of team cognition: A process-oriented theory of knowledge emergence in teams. *Journal of Applied Psychology*, 101(10), 1353-1385.
- Gupta, A. K., Smith, K. G., & Shalley, C. E. (2006). The interplay between exploration and exploitation. *Academy of Management Journal*, 49(4), 693-706.
- Hackman, J. R. & Morris, C. G. (1975). Group tasks, group interaction processes, and group performance effectiveness: A review and proposed integration. *Advances in Experimental Social Psychology*, 8, 45-99.
- Harrison, J. R., & Carroll, G. R. (1991). Keeping the faith: A model of cultural transmission in formal organizations. *Administrative Science Quarterly*, 552-582.
- Harrison, J. R., Lin, Z., Carroll, G. R., & Carley, K. M. (2007). Simulation modeling in organizational and management research. *Academy of Management Review*, 32(4), 1229-1245.
- Hart, S. L. (1992). An integrative framework for strategy-making processes. *Academy of Management Review*, 17(2), 327-351.
- Hedberg, B. (1981). How organizations learn and unlearn? In P. C. Nystrom and W. H. Starbuck,

- (Eds.) *Handbook of organizational design*. London: Oxford University Press, 8-27.
- Henrich, J., & Boyd, R. (1998). The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behavior*, 19(4), 215-241.
- Henrich, J. (2001). Cultural transmission and the diffusion of innovations: Adoption dynamics indicate that biased cultural transmission is the predominate force in behavioral change. *American Anthropologist*, 103(4), 992-1013.
- Huber, G. P. (1991). Organizational learning: The contributing processes and the literatures. *Organization Science*, 2(1), 88-115.
- Hulin, C. L., & Ilgen, D. R. (2000). Introduction to computational modeling in organizations: The good that modeling does. In D.R. Ilgen & C.L. Hulin (Eds.) *Computational Modeling of Behavior in Organizations*. American Psychological Association. Washington, D.C.
- Jehn, K. A., Northcraft, G. B., & Neale, M. A. (1999). Why differences make a difference: A field study of diversity, conflict and performance in workgroups. *Administrative Science Quarterly*, 44(4), 741-763.
- Katila, R., & Ahuja, G. (2002). Something old, something new: A longitudinal study of search behavior and new product introduction. *Academy of Management Journal*, 45(6), 1183-1194.
- Kontoghiorghes, C., Awbre, S. A., & Feurig, P. L. (2005). Examining the relationship between learning organization characteristics and change adaptation, innovation, and organizational performance. *Human Resource Development Quarterly*, 16(2), 185-212.
- Kotter, J. (2012). How the most innovative companies capitalize on today's rapid-fire strategic challenges-and still make their numbers. *Harvard Business Review*, 90(11), 43-58.
- Kozlowski, S. W., & Ilgen, D. R. (2006). Enhancing the effectiveness of work groups and teams. *Psychological Science in the Public Interest*, 7(3), 77-124.
- Kozlowski, S. W. J., & Klein, K. J. (2000). A multilevel approach to theory and research in organizations: Contextual, temporal, and emergent processes. In K. J. Klein & S. W. J. Kozlowski (Eds.), *Multilevel theory, research, and methods in organizations: Foundations, extensions, and new directions* (pp. 3-90). San Francisco: Jossey-Bass.
- Lau, D. C., & Murnighan, J. K. (1998). Demographic diversity and faultlines: The compositional dynamics of organizational groups. *Academy of Management Review*, 23(2), 325-340.
- Levinthal, D. A., & March, J. G. (1993). The myopia of learning. *Strategic Management Journal*, 14(S2), 95-112.



- Lipson, H. (2007). Principles of modularity, regularity, and hierarchy for scalable systems. *Journal of Biological Physics and Chemistry*, 7(4), 125.
- Mahajan, V., Muller, E., & Bass, F. M. (1995). Diffusion of new products: Empirical generalizations and managerial uses. *Marketing Science*, 14(3), G79-G88.
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71-87.
- March, J. G., Schulz, M., & Zhou, X. (2000). *The dynamic of rules: Change in written organizational codes*. Stanford, CA: Stanford University Press.
- Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., ... & Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, 2(3), 191-215.
- Miles, R. E., Snow, C. C., Meyer, A. D., & Coleman Jr., H. J. (1978). Organizational strategy, structure, and process. *Academy of Management Review*, 3(3), 546-562.
- Miller, K. D., Zhao, M., & Calantone, R. J. (2006). Adding interpersonal learning and tacit knowledge to March's exploration-exploitation model. *Academy of Management Journal*, 49(4), 709-722.
- Miner, A. S., & Mezias, S. J. (1996). Ugly duckling no more: Pasts and futures of organizational learning research. *Organization Science*, 7(1), 88-99.
- Moore, G. E. (1965). Cramming more components onto integrated circuits. *Electronics Magazine*.
- O'Reilly, C. A., & Tushman, M. L. (2004). The ambidextrous organization. *Harvard Business Review*, 82(4), 74-83.
- Ostroff, C., & Kozlowski, S. W. (1992). Organizational socialization as a learning process: The role of information acquisition. *Personnel Psychology*, 45(4), 849-874.
- Page-Jones, M. (1980). *The practical guide to structured systems design*. New York: Yourdon Press.
- Pelled, L. H., Eisenhardt, K. M., & Xin, K. R. (1999). Exploring the black box: An analysis of work group diversity, conflict and performance. *Administrative Science Quarterly*, 44(1), 1-28.
- Perretti, F., & Negro, G. (2006). Filling empty seats: How status and organizational hierarchies affect exploration versus exploitation in team design. *Academy of Management Journal*, 49(4), 759-777.

- Polanyi, M. (1967). *The tacit dimension*. Garden City, NY: Anchor Books.
- Popper, K. R. (1972). *Objective knowledge: An evolutionary approach*. Clarendon Press.
- Rodan, S. (2005). Exploration and exploitation revisited: Extending March's model of mutual learning. *Scandinavian Journal of Management*, 21(4), 407-428.
- Rodriguez, D., Sicilia, M. A., Garcia, E., & Harrison, R. (2012). Empirical findings on team size and productivity in software development. *Journal of Systems and Software*, 85(3), 562-570.
- Schilling, M. A. (2000). Toward a general modular systems theory and its application to interfirm product modularity. *Academy of Management Review*, 25(2), 312-334.
- Schneider, B. (1987). The people make the place. *Personnel Psychology*, 40(3), 437-453.
- Senge, P. M. (1995). *The fifth discipline: The art and practice of the learning organization*. Crown Pub.
- Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6), 467-482.
- Simon, H. A. (1991). Bounded rationality and organizational learning. *Organization Science*, 2(1), 125-134.
- Stasser, G., & Titus, W. (1985). Pooling of unshared information in group decision making: Biased information sampling during discussion. *Journal of Personality and Social Psychology*, 48, 1467-1478.
- Staw, B. M. (1980). The consequences of turnover. *Journal of Occupational Behavior*, 1(4), 253-273.
- Steiner, I. D. (1972). *Group processes and group productivity*. New York: Academic.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge: MIT Press.
- Taylor, A., & Greve, H. R. (2006). Superman or the fantastic four? Knowledge combination and experience in innovative teams. *Academy of Management Journal*, 49(4), 723-740.
- United States Census Bureau (2015). 2015 SUSB annual datasets by establishment industry. Retrieved March 24, 2018 from <https://www.census.gov/data/datasets/2015/econ/susb/2015-susb.html>.
- Valve Corporation (2012). *Handbook for new employees: A fearless adventure in knowing what to do when no one's there telling you what to do*. Bellevue, WA: Valve Press.

- Watts, D. J. (1999). Networks, dynamics, and the small-world phenomenon. *American Journal of Sociology*, 105(2), 493-527.
- Wegner, D. M. (1987). Transactive memory: A contemporary analysis of the group mind. In *Theories of group behavior* (pp. 185-208). New York: Springer.
- Weick, K. E. (1976). Educational organizations as loosely coupled systems. *Administrative Science Quarterly*, 1-19.
- Weinhardt, J. M., & Vancouver, J. B. (2012). Computational models and organizational psychology: Opportunities abound. *Organizational Psychology Review*, 2(4), 267-292.
- Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding, and selection in evolution. *Proceedings of the VI International Congress of Genetics*, 1, 356-366.
- Zohar, D., & Hofmann, D. A. (2012). Organizational culture and climate. In S. W. J. Kozlowski (Ed.), *The Oxford handbook of industrial and organizational psychology* (pp. 643-666). Oxford, UK: Oxford University Press.