

**Transactive Systems Model of Collective Intelligence:
The Emergence and Regulation of Collective Attention, Memory, and Reasoning.**

by
Pranav Gupta

A dissertation submitted in partial fulfillment of the requirements for
the degree of Doctor of Philosophy (Industrial Administration, Organizational
Behavior and Theory) to the Tepper School of Business at Carnegie Mellon University

2022

Doctoral Committee

Anita Williams Woolley (Chair), Carnegie Mellon University

Kathleen M. Carley, Carnegie Mellon University

Oliver Hahl, Carnegie Mellon University

Melissa Valentine, Stanford University

Copyright © 2022 by Pranav Gupta

All Rights Reserved

To my infinitely supportive family—

Ved, Shivam, Pragya, Mom, and Dad

Abstract

Modern organizations face a highly complex and dynamic environment. In this dissertation, I argue that collectives (teams, organizations, or communities) facing highly complex and dynamic situations need to be designed for collective intelligence, or the ability to achieve goals in a wide range of environments. Integrating and building on extant work, I theorize a *Transactive Systems Model of Collective Intelligence*, guided by agent-based modeling, and test it with data from open-source software teams. I also explore and propose a set of diagnostic indicators as an extension for further development, to enable monitoring and intervention by algorithmic tools and/or leaders.

I take a complex adaptive systems view of collectives and describe how transactive attention, memory, and reasoning systems (TAS, TMS, and TRS) emerge from individual-level cognitive processes and member interactions to shape the emergence of collective intelligence. I further theorize how these systems interact with each other and respond to dynamic environmental complexity. I complement this narrative theory with agent-based modeling (ABM) to validate the sufficiency of the proposed transactive process. Once validated, I use the ABM to conduct two virtual experiments to demonstrate the co-regulation of TMS and TAS, which is complemented by TRS. Based on these virtual experiments, I derive hypotheses about the critical environmental threats that each transactive system is specifically equipped to address, highlighting the co-regulation necessary in response to changes in the environment which we theorize underlie the development and maintenance of collective intelligence. Finally, I empirically test corresponding hypotheses for aggregated system behavior by analyzing 18 months of archival data from 476 open-source software teams. Consistent with predictions, I find evidence confirming the hypotheses and providing initial support for the transactive systems theory.

In the next chapter, I build on the socio-cognitive architecture of collective intelligence articulated in the Transactive Systems Model and theorize three observable collaborative processes that are related to the transactive system processes. I propose that these observable processes can serve as diagnostic indicators to provide real-time information about the functioning of the underlying, largely unobservable complex adaptive system. I explore these collaborative process indicators in another virtual experiment that supports their general utility in signaling the level of functioning of different transactive systems. I also propose combining them into a single metric that acts as a leading indicator of collective intelligence. It can be monitored in real-time and guide the diagnosis of underlying problems by looking at its component parts. Articulation and refinement of such metrics provide useful guides for intervention by humans or algorithmic tools and together with transactive systems lay a foundation for a Machine Theory of Collective Intelligence.

While much remains to be learned about the nature of collective intelligence, this dissertation presents a multi-method systems approach for investigating its emergence, adaptation, and diagnosis, laying the groundwork for future research.

Keywords: collective intelligence, team cognition, transactive attention system, transactive memory system, complex adaptive systems, agent-based simulations, machine theory of collective intelligence

Acknowledgements

I am in debt to the faculty members and colleagues at Carnegie Mellon University for providing a rich, interdisciplinary, and stimulating intellectual environment for my doctoral work, but in particular to Anita Woolley. This work was shaped under her guidance and would not have been possible without her nurturing support. Her flexible approach has brought scientific rigor to my imagination. The years spent intellectually sparring with her over these ideas have made me a better researcher and I, happily, look forward to many more to come. To Kathleen Carley, I owe a special debt. She has helped me look at social science, systematically, from a modeling perspective and built my computational research muscles. Your insights and critiques have always cut to the heart of the ideas and resulted in conceptual clarity and simplicity. I endeavor to emulate this skill. I give my thanks to Melissa Valentine, Oliver Hahl, Linda Argote, Ethan Bernstein, and Anna Mayo for reading, critiquing, and providing useful feedback on earlier versions of this manuscript. It has improved the quality of ideas and spurred exciting new directions.

All of this work was made possible because of a number of grant fundings and awards from the Tepper School, DARPA, ARO, and NSF. I continue to appreciate the support and confidence shown by various anonymous reviewers and committee members who found these ideas exciting.

To the late Herbert Simon and Vivekananda, I owe the ultimate debt. First, their papers and essays exploring the nature of human cognition and its implications for organizational and societal design have exerted a major influence on my thinking about management research and continue to serve as an inspiration. Second, their approach to the design of educational systems encourages me to be a better teacher.

I have been fortunate to have found good friends and supportive colleagues along the way. This is an expression of my gratitude to all of you for conversations on research or otherwise - Varun Deshpande, Ella Glikson, Fade Eadeh, Young Ji Kim, Jonathan Kush, Anna Mayo, Jerry Guo, Alessandro Iorio, Elizabeth Campbell, Rafael Kaufman, Heather Altman, Eric Silver, KiWon Haan, Ben Ostrowski, Sae Park, Jisoo Park, and members of the Herbsleb Lab. Special mention to Marat Vailev for his super-human help with data collection used in this work as well as Lawrence Rapp and Laila Lee for making sure I was well-fed, well-informed, and all around being excellent at tackling all administrative situations I found myself in.

Finally, what can I even say about my family —Mom, Dad, Shivam, Pragya, and Ved— to whom I dedicate this work. I am lucky to have you all in my corner. As always, it has been a wild ride that I would happily take again with all of you. To many more adventures that are equal parts scary and exciting!

Table of Contents

Acknowledgements	v
Table of Contents	vi
List of Tables	viii
List of Figures	ix

Introduction:

Limits of Structural Control and Organizing in the Face of Continuous Change	1
---	----------

Transactive Systems Model of Collective Intelligence:

The Emergence and Regulation of Collective Attention, Memory, and Reasoning.	5
---	----------

Theorizing Collective Intelligence in Complex Adaptive Systems	7
Emergence of Transactive Memory, Attention, and Reasoning Systems	10
Emergence of Transactive Memory System (TMS)	11
Emergence of Transactive Attention System (TAS)	16
Emergence of Transactive Reasoning System (TRS)	20
Emergence of Collective Intelligence via Adaptive Regulation of Transactive Systems	24
Efficiency Function: Co-Regulation of TMS and TAS for Efficient Resource Use	26
Maintenance Function: Regulate TRS for Effectively Exploring & Retaining Resources	29
Collective Intelligence in Open Source Software Teams:	
An Empirical Investigation of Transactive Systems Model	31
Overview of Research Setting: Open Source Software (OSS) Project Teams	32
Data and Operationalization of Constructs	33
Results	37
General Discussion	40
Implications	42
Limitations And Future Directions	43
Conclusion	45

Dynamic Indicators of Collective Intelligence:

The Diagnostic Role of Collaborative Processes in Transactive Systems Model	46
--	-----------

Four Common Problems That Plague the Management of Complex Adaptive Systems	47
Collaborative Processes and their Relationship with Transactive Systems Model	50
Sufficient Application of Collective Effort	53

Appropriate Use of Performance-Strategies	56
Sufficient Application of Knowledge and Skill	57
Identify Leverage Points by Putting All of Them Together	59
Exploring Diagnostic Use of Collaborative Indicators via Simulation	60
Model overview and Virtual Experiment Design	61
Indicators Operationalization	62
Examining Results and Exploring the Diagnostic Role	63
General Discussion	67
Implication for Leadership and Coaching Functions	68
Foundations for the Development of AI-based Team Facilitator or Coach	69
Conclusion	73
Summary Conclusions and Future Directions	74
Future Exploration: Developing Organizational Design Principles and Digital Tools for Minimum Viable Collective Intelligence	77
Managerial Implications	82
Conclusion	85
Appendix	86
References	88

List of Tables

Table 2.1. <i>OSS-GitHub Study Descriptive Statistics and Correlations</i>	39
Table 2.2. <i>Test of Hypotheses in OSS-GitHub: OLS Regression Results</i>	39
Table A1. <i>Procedure for Computational Model for Emergence of Collective Intelligence.</i>	86
Table A2. <i>Rulesets for Triaging a Member's Inbox.</i>	87

List of Figures

- Figure 1.1.** *Environmental dynamism and members' need for higher autonomy puts a limit on our ability to use structural controls to organize efficiently.* 3
- Figure 2.1.** *Collective intelligence emerges as a result of adaptive regulation of three socio-cognitive systems (transactive memory, attention, and reasoning systems) in response to their changing environmental complexity.* 9
- Figure 2.2.** *The socio-cognitive architecture of individual- to collective-level emergence of transactive memory, attention, and reasoning systems (TMS, TAS, and TRS).* 12
- Figure 2.3.** *Model validation for the emergence of (A) TMS, maintains high specialization over time; (B) TAS, maintains high convergence over time; and (C) TRS, maintains high commitment over time.* 15
- Figure 2.4.** *Overview of adaptive regulation of transactive systems. TMS and TAS regulate each other for efficient utilization of resources (see inside left panel). While TRS regulates the tradeoff between maintaining resources and seeking efficiency (balance resources across panels).* 25
- Figure 2.5.** *Virtual Experiment Results: (Top two rows) Efficiency Function: Collectives with co-regulated TMS and TAS outperforms TMS-only and TAS-only Collectives across all conditions. (Bottom row) Maintenance Function: When member effort is at risk, collectives with TMS, TAS, and TRS exhibit higher intelligence than collectives with TMS and TAS only.* 28
- Figure 2.6.** *High workload projects without a well-developed TAS suffer the most. Collective intelligence is operationalized as NPM score - an index capturing the projects' ability to maintain good quality code.* 40
- Figure 3.1.** *Collaborative processes dynamically sense the changing environmental complexity (indicator of current collective intelligence) and help the collective accurately diagnose the problem and allow managers to take actions that desirably regulate transactive systems (calibrated response).* 51
- Figure 3.2.** *Collaborative processes can be used to triangulate and diagnose possible problems with different transactive processes, the underlying mechanism for collective intelligence. Taken together, the collaborative processes indicate the current level of collective intelligence.* 54
- Figure 3.3.** *Virtual Experiment Results: Composite, process-based indicator of collective intelligence (A) along with three underlying collaborative process indicators (B). The performance-based measure of collective intelligence (C) along with underlying performance score over time (D).* 65
- Figure 3.4.** *Potential roles for AI-based team facilitator or coaching agents for improving collective intelligence. Adapted from (Gupta and Woolley, 2021)* 70
- Figure 4.1.** *Organizational forms located in the Efficiency x Maintenance Graph.* 78

I

Introduction:

Limits of Structural Control and Organizing in the Face of Continuous Change

Modern work has shifted towards increasingly dynamic, temporary, and decentralized forms of organizing which is achieved by increasing the permeability of organizational boundaries. Permeable boundaries enable the breaking of structural silos which allows access to valuable resources across areas of the organization. Consequently, there has been a sharp increase of organizing activities reflecting this boundary permeability such as multi-teaming (Cummings & Haas, 2012; O’Leary et al., 2011), dynamic teams (Mayo, in press), role fluidity, overlapping work, and worker dispersion (Bakker, 2010; Mortensen, 2014; O’Leary & Cummings, 2007; Saunders & Ahuja, 2006; Valentine & Edmondson, 2015) facilitated by organizing technology such as new channels, platforms, and devices (Bernstein et al., 2016; Lee & Edmondson, 2017; Leonardi & Vaast, 2017; Mortensen & Haas, 2018; Neeley & Leonardi, 2018). In some settings, these practices have led to more productivity (Ancona & Caldwell, 1992; Faraj & Yan, 2009; Hackman & Wageman, 1995), more innovation (Dougherty & Dunne, 2012), and more collaboration (Bechky, 2003; Kellogg et al., 2006). This represents a significant shift from the central assumptions of the classical administrative theory that values boundaries as a necessary form of structural control needed for efficiently retaining and organizing bounded rational employees (March & Simon, 1993, Simon, 1997).

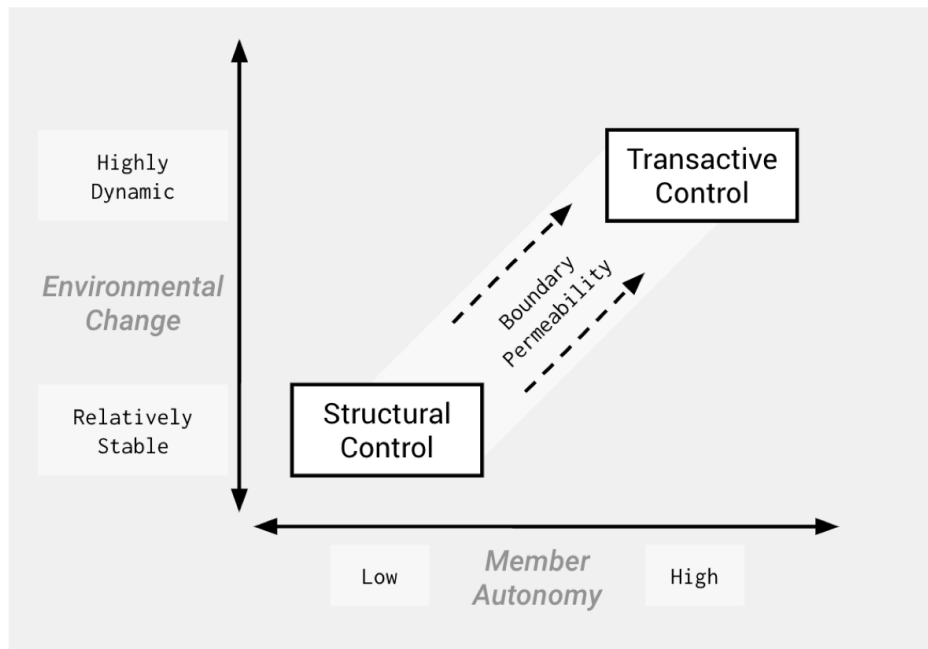
However, the management literature has not been able to offer much in terms of a systematic path forward as it becomes clear that traditional organizational structures are incompatible with the demands of complex work environments. The loosening or disruption of structural control mechanisms to enable highly dynamic work are also problematic. It results in errors and inefficiencies due to the increased likelihood of worker overload and burnout. There have been increasing calls for management theory to adapt its assumptions

about organizations to more appropriately reflect their reality as complex adaptive systems (Arrow et al., 2000; Cronin et al., 2011; Hackman, 2012; Kozlowski & Klein, 2000; Mathieu et al., 2019). This has resulted in the growing use of concepts originating in research on intelligence in humans and machines (Csaszar & Steinberger, 2021; March, 2006; Simon, 1998), which has for decades studied the functions that enable systems to adapt and accomplish goals in a wide range of environments that vary in complexity (Legg & Hutter, 2007). And as we see the rapid development of artificial intelligence and its integration into all areas of work, it seems increasingly appropriate to view organizations less as stable structures and more as complex adaptive systems, designed to integrate human intelligence and machine intelligence into collective intelligence. My goal is to embrace the complexity, to better understand the underlying dynamics in order to facilitate the emergence of desirable properties.

In this dissertation, I introduce a *Transactive Systems Model of Collective Intelligence*, which provides a process-based framework for understanding how coordination can be accomplished in a complex adaptive system. In a transactive system, members share and allocate cognitive resources and coordinate collective action to their mutual benefit. Instead of using *structural control* to organize and retain employees by administratively bounding organizational situations, this theory proposes that autonomy-seeking members can exercise *transactive control* to engage each other's bounded cognitive capacities in a highly dynamic environment (see Figure 1.1). Specifically, I theorize a socio-cognitive architecture that assumes cognitive agents—humans or machines—use their metacognition (understanding of their own and others' memory, attention, and reasoning processes) to transactively influence each other's individual decision-making in order to negotiate and maintain mutually rewarding goals and efficiently coordinate collective action in a continually changing environment. The work unfolds via agent-based models to support and

Figure 1.1

Environmental dynamism and members' need for higher autonomy puts a limit on our ability to use structural controls to organize efficiently.



develop the associated narrative theory and paired with analysis of archival data to examine similarities and differences between computational models and patterns in the real world.

In Chapter 2 (work conducted with Anita Woolley and Kathleen Carley), I take a complex adaptive systems view of collectives and describe how transactive attention, memory, and reasoning systems (TAS, TMS, and TRS) *emerge* from individual-level cognitive processes and member interactions to shape the emergence of collective intelligence. I further theorize how these systems *adaptively regulate* each other and respond to dynamic environmental complexity. I complement this narrative theory with agent-based modeling (ABM) to validate the sufficiency of the proposed transactive processes. Once validated, I use the ABM to conduct two virtual experiments to demonstrate the co-regulation of TMS and TAS, which is complemented by TRS. Based on these virtual experiments, I derive hypotheses about the critical environmental threats that each transactive system is

specifically equipped to address, highlighting the co-regulation necessary in response to changes in the environment which we theorize underlie the development and maintenance of collective intelligence. Finally, I test these hypotheses by analyzing 18 months of archival data from 476 open-source software teams. Consistent with predictions, I find evidence confirming the hypotheses and providing initial support for the transactive systems theory.

In chapter 3 (work with Anita Woolley), I build on the socio-cognitive architecture of collective intelligence and theorize three observable collaborative processes— level of collective effort, use of performance strategy, and application of knowledge and skill— that are related to the ongoing regulatory states of the transactive systems. I propose that these observable processes can serve as diagnostic indicators to provide real-time information about the *non-linear* and *path-dependent* behaviors of the underlying, largely unobservable complex adaptive system that is prone to four common types of problems. I explored these collaborative process indicators in an agent-based simulation and investigated their general utility in signaling the level of regulation achieved by different transactive systems. I also examine their validity as a real-time leading indicator of collective intelligence by combining the individual process indicators and comparing it to the performance-based assessment of collective intelligence. The real-time indicator offers the benefit of enabling monitoring and diagnosis of underlying problems along with pointers to the transactive systems potentially involved in difficulties. Articulation and refinement of such metrics hold the potential to provide useful guides for intervention by humans or algorithmic tools and together with transactive systems lay a foundation for a Machine Theory of Collective Intelligence.

In the final chapter, I conclude by summarizing and considering the theoretical and empirical work as a whole and offering some thoughts on new forms of decentralized organizations and designing tools for autonomously improving collective intelligence.

II

Transactive Systems Model of Collective Intelligence:

The Emergence and Regulation of Collective Attention, Memory, and Reasoning.

With globalism and technology-enabled hyper-connectedness, the context surrounding organizations has become more complex and dynamic. Consequently, collective work is characterized by increasing uncertainty and fluidity (Edmondson and Harvey, 2018; Eisenman et al., 2020) creating challenges for the traditional organizational structures and conditions previously viewed as foundational to enabling collectively-rational behavior in boundedly rational actors (March et al., 1993). In response to this growing dynamism, there have been increasing calls for management theory to reflect the reality of collective work operating less like a “machine” and more like a complex adaptive system (Arrow et al., 2000; Cronin et al., 2011; Hackman, 2012; Kozlowski & Klein, 2000; Mathieu et al., 2019). This is partially reflected in the management literature by the growing use of concepts originating in research on intelligence in humans and machines (Csaszar & Steinberger, 2021; March, 2006), which has for decades studied the functions that enable systems to adapt and accomplish goals in a wide range of environments that vary in complexity (Legg & Hutter, 2007). And as we see the rapid development of artificial intelligence and its integration into all areas of work, it seems increasingly appropriate to view organizations less as stable structures and more as complex adaptive systems, designed to integrate human intelligence and machine intelligence into collective intelligence.

However, while a growing body of work points to a variety of correlates and predictors of collective intelligence (Bernstein et al. 2018; Leonard & Levin, 2022; Riedl et al., 2021; Woolley et al., 2010), we lack a detailed process theory for how to design a system for collective intelligence. Integrating and building on extant work, here we articulate a Transactive Systems Model of Collective Intelligence. Taking a complex adaptive systems

view of collectives, we propose systems theory articulating how transactive memory (TMS), attention (TAS), and reasoning (TRS) systems emerge from individual-level cognitive processes and member interactions to shape the emergence of collective intelligence. We further theorize how these systems interact with each other and dynamically respond to changing environmental complexities. We describe the theory both narratively as well as formally by operationalizing it in an agent-based model, which we validate against expected emergent patterns. We extend the simulation and use it in virtual experiments to develop hypotheses about the emergence and adaptation of collective intelligence in contexts that vary in environmental complexity. The theory and simulation inform hypotheses that we then test empirically in an archival dataset from 476 open-source software teams.

Taken together, our theory, simulation, and empirical analysis make three key contributions to the literature. First, we contribute to the growing literature on collective intelligence and social cognition by integrating and building on extant work to articulate a process theory of collective intelligence and the underlying transactive systems that drive the emergence and adaptation of collective memory, attention, and reasoning functions. Second, we formalize this narrative theory in a computational model, which enables us to validate the sufficiency of theorized mechanisms underlying the emergence of collective memory, attention, and reasoning and demonstrate how their adaptive interactions result in collective intelligence. And, finally, we operationalize constructs for use in analyses of archival data to enable an empirical test of the behavioral patterns we theorize we can observe as a result of these underlying processes in real-world environments. This initial empirical test provides support for aggregate behavior emerging from underlying dynamics that are theorized and calls for future research in other settings where the dynamics can be tested directly. Thus, this work which theorizes a socio-cognitive architecture for collective intelligence represents an important foundation for future research and theory about organizational functioning.

Theorizing Collective Intelligence in Complex Adaptive Systems

Intelligence, whether in the context of an individual, a collective, or a technological system, is broadly defined as the ability to perform or achieve goals in a wide range of environments (Legg & Hutter, 2007). Implicit in “ability to perform” is the notion of adapting in order to maintain performance, and “wide range of environments” refers to varying complexity of the task and resource environment in which the collective operates. Here, complexity captures the dominant environmental threats in terms of uncertainty in task rewards and member resources, level of task workload, and knowledge interdependencies at a given time (Raveendran et al., 2020). Thus, high collective intelligence is signaled by sustained performance in the face of changes in complexity over a given time period:

$$Collective\ Intelligence = \sum_{time\ period} Complexity_{time}^{task} * Performance_{task}^{collective} \dots eq\ (1)$$

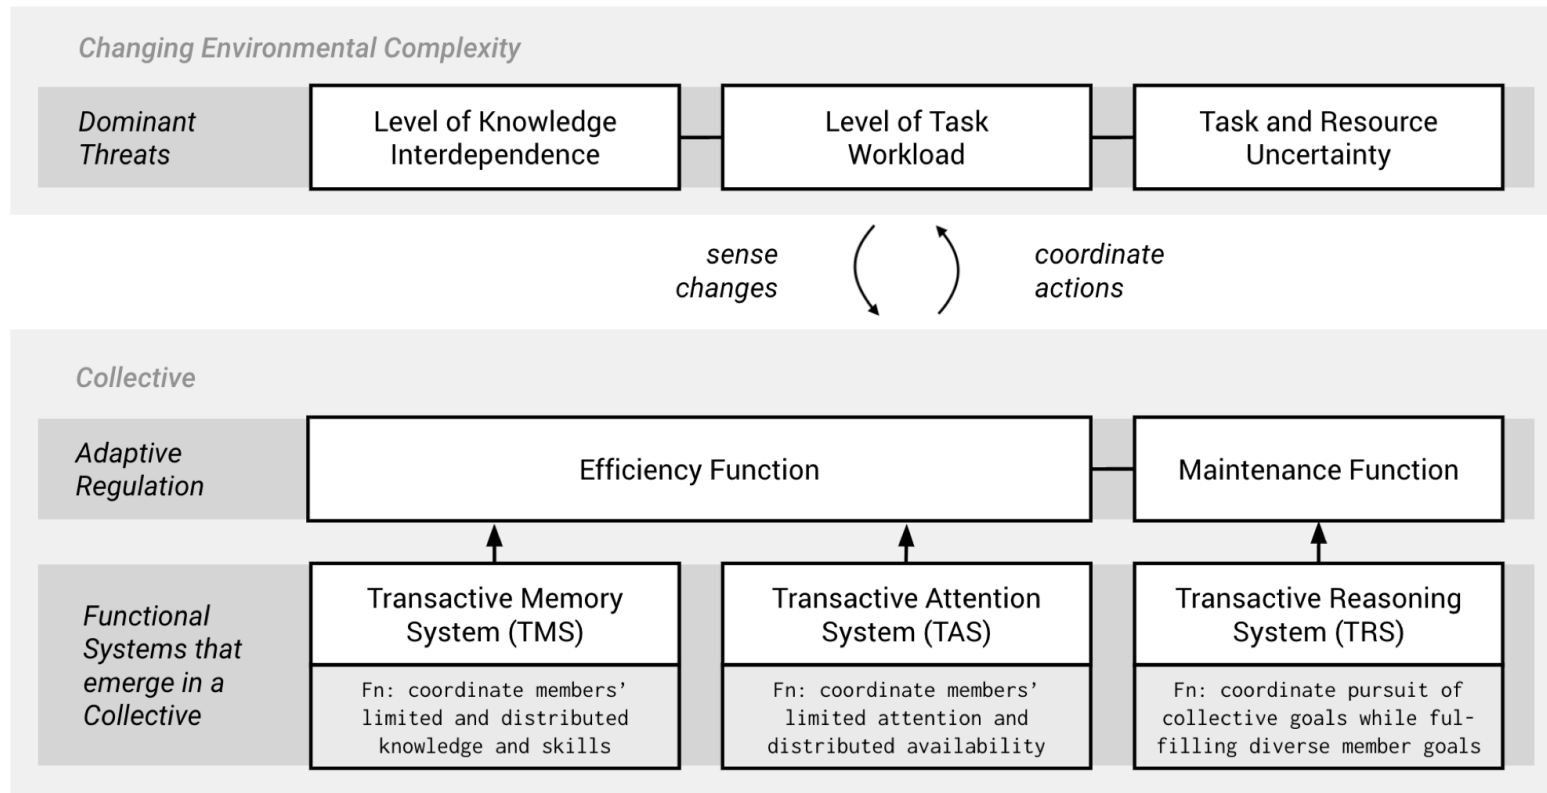
Extant work suggests that three core cognitive functions underlying intelligence in any system are memory, attention, and reasoning processes. These functions are essential to intelligent functioning whether the system is the human brain (Luria, 1973) or another biological, technological, or hybrid system (Malone & Bernstein, 2015). Thus, a foundation for the emergence of intelligence is the development of systems accomplishing these functions. In line with this, we propose that three functionally distinct socio-cognitive systems emerge in collectives: (a) Transactive memory system (TMS) for governing the coordination of members’ limited and distributed knowledge and skills; (b) Transactive attention system (TAS) for governing the coordination of members’ limited attention and distributed availabilities; and (c) Transactive reasoning system (TRS) for governing the coordination of effective collective goals while fulfilling members’ diverse motivations. The emergence and mutual regulation of these systems provide the foundation for the emergence of collective intelligence, enabling the collective to adapt in response to environmental

complexity. The mutual adaptation of these systems is central to the collective's ability to accomplish two generic functions long thought to be essential in all social systems (Arrow et al., 2000; Hackman, 1987; Taylor, 2012): first, to accomplish the system's goals by efficiently utilizing its resources (i.e. efficiency function); and, second, to maintain the system and its resources by pursuing appropriate goals while fulfilling members' needs (i.e. maintenance function). The relative importance of these two functions is also dynamic, differs by context, and changes over time. Thus, the design problem of collective intelligence is to put together a set of heterogeneous, boundedly rational members who can coordinate their distributed cognitive resources and diverse personal goals to formulate a series of joint decisions and actions that fulfill its efficiency and maintenance functions over time (Figure 2.1).

In articulating a theory describing complex adaptive systems that exhibit a phenomenon like collective intelligence, it is important to attend to both the mechanisms for dynamics of emergence and adaptation. Emergent or bottom-up systems dynamics involve describing the nature of lower-level actors and the rules that guide their interactions with others to theorize recognizable patterns of behavior exhibited at the higher level over time. A common example for illustration is bird flocking behavior, where the higher-level behavior (flocking) is described by the rules governing how the individual birds decide where to position themselves (e.g., keep X distance from each other). A goal in articulating these rules underlying the emergence of a phenomenon is not to be exhaustive but rather to determine the minimum rule specification necessary to produce the pattern, also referred to as "generative sufficiency" (Epstein, 1999; Kozlowski et al., 2013; Kozlowski & Klein, 2000). Theorizing about adaptation, by contrast, typically uses a top-down perspective and involves specifying how the different elements of the system form feedback loops with each other and differentially respond to different environmental variables. Adaptive behavior often arises as the relative strengths of competing feedback loops shift as the system responds to changes in

Figure 2.1

Collective intelligence emerges as a result of adaptive regulation of three socio-cognitive systems (transactive memory, attention, and reasoning systems) in response to their changing environmental complexity.



the environment (Meadows & Wright, 2008). For example, a thermostat forms a feedback loop with the room temperature and automatically controls the heating, ventilation, and air conditioning systems to maintain the desired temperature. Explicitly theorizing about both emergence and adaptation enables us to identify the mechanisms that generate, sustain, and conditionally modify relevant systems behaviors in response to environmental variables.

Following this approach, we first describe the processes that lead to the emergence of three transactive systems (TMS, TAS, and TRS), and then articulate the adaptive regulation of the transactive systems that provide the foundation of collective intelligence. As we do so, we first articulate our theory in narrative form, and then operationalize it in an agent-based computational model, enabling us to validate that the rules implied by our theory result in predicted patterns in simulated data and conduct virtual experiments (Carley, 2001). The simplified model descriptions and simulation outcomes are interwoven with our theoretical arguments (See appendix for a summary of procedure and rulesets).

Emergence of Transactive Memory, Attention, and Reasoning Systems

In this section, we articulate the individual cognitive and inter-member processes that lead to the emergence of transactive memory (TMS), attention (TAS), and reasoning (TRS) systems (see Figure 2.2 for an overview). For each transactive system, we start by theorizing the individual cognitive, meta-cognitive, and inter-member processes. We then instantiate these processes in a baseline agent-based model and observe collective-level patterns of behavior that emerge, thereby allowing us to validate their generative sufficiency. In the next section, we progressively combine these baseline models in “virtual experiments” to explore how their interaction under different environmental threats exhibits adaptive behavior. The results of the virtual experiments inform our hypotheses, which we then test against our empirical data in the section after.

Emergence of Transactive Memory System (TMS)

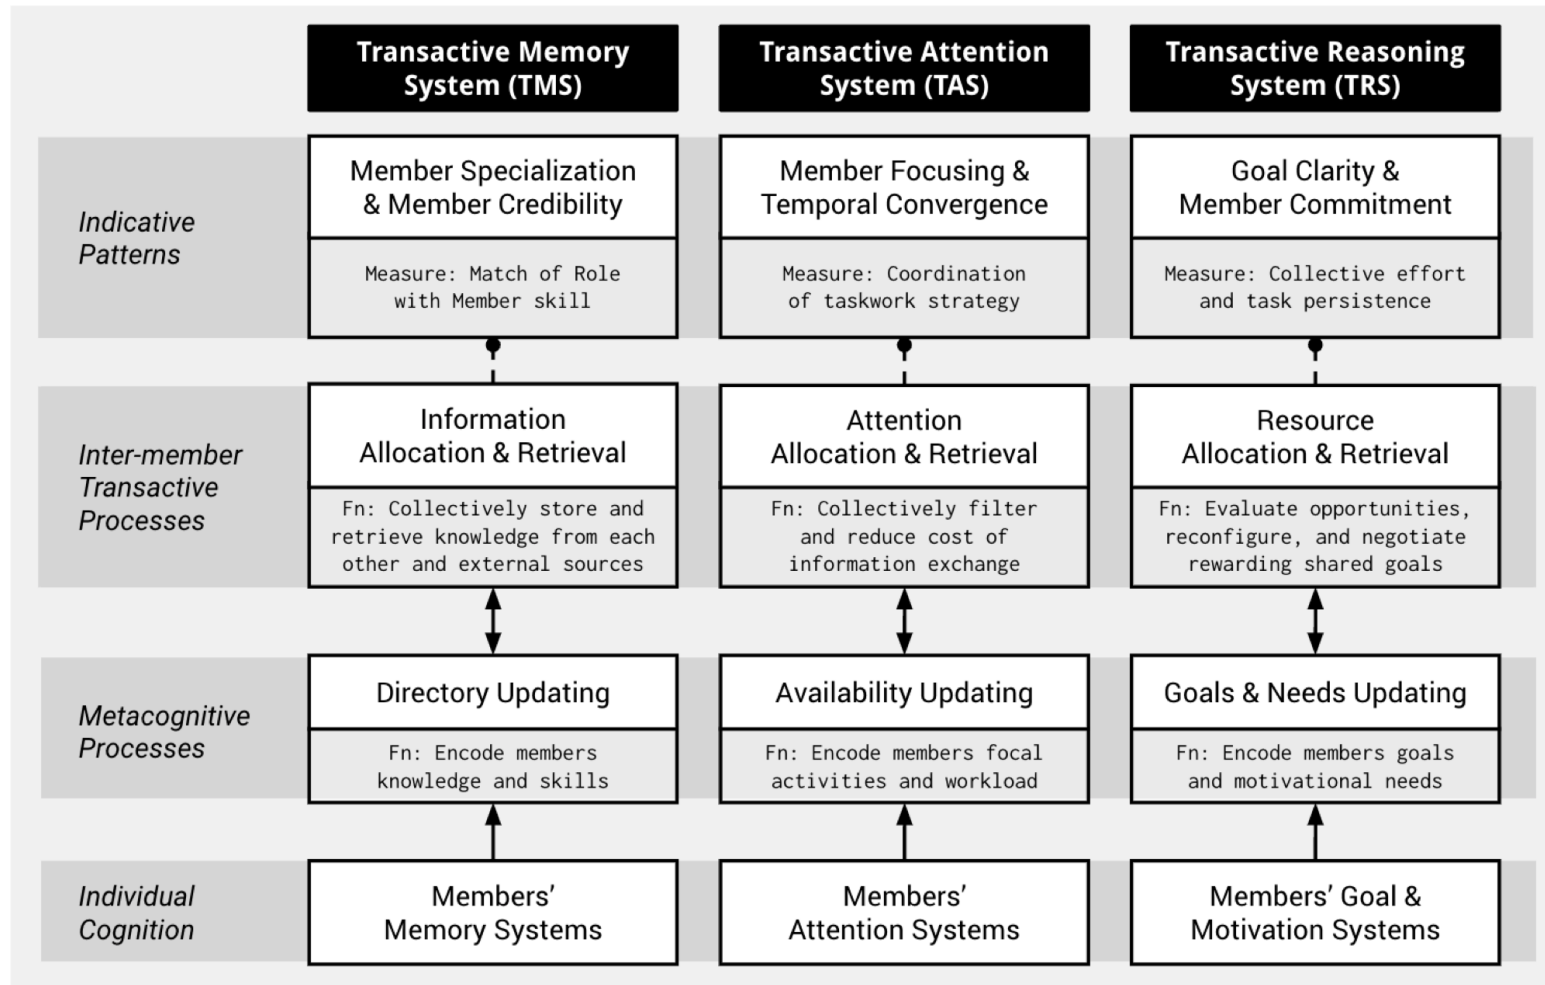
Collective memory is an essential element in determining the amount of knowledge and expertise that is available to be applied to collective work. A transactive memory system (TMS) is a form of individually-held cognitive understanding of team members' skills, whereby each member is aware of who knows what on the team and relies on retrieving that knowledge to achieve effective coordination (Ren & Argote, 2011). A well-functioning TMS greatly improves how collective memory is organized and maintained. Work by Wegner (1987, 1995) conceptualized TMS as a dynamic system comprising of (1) members' current level of knowledge and skills (i.e. individual memory system); (2) their current understanding of each others' task-relevant knowledge and skills (i.e. individual meta-memory); and (3) the inter-member processes that facilitate the dynamic updating of meta-memory and subsequent allocation and retrieval of knowledge to and from the most appropriate member (Figure 2.2).

Individual Memory and Meta-Memory

Individual memory systems have a limited capacity to hold knowledge, which is managed via four memory processes: encoding, storage, retrieval, and forgetting (Newell & Simon, 1972). During the encoding process, the incoming information is designated a relevant "label." This is followed by the storage process wherein the information is assimilated in memory along with knowledge with similar labels. The labels are used to guide the retrieval process whenever a specific piece of knowledge is recalled. Finally, forgetting happens automatically whereby knowledge that is not accessed for a long time starts to decay. Alongside these processes, individuals also develop meta-memory or a cognitive directory of all the labels and their location in the person's memory. The contents of meta-memory are very influential in determining the speed of memory processes. Meta-memory can also serve as a directory for labels that are not stored in a person's own memory, but in an external source (e.g. another person's memory system, or a specific physical or digital archive; Flavell

Figure 2.2

The socio-cognitive architecture of individual- to collective-level emergence of transactive memory, attention, and reasoning systems (TMS, TAS, and TRS).



& Wellman, 1977). Having this directory drastically reduces the time needed to look up the location of relevant knowledge.

Inter-Member TMS Processes

TMS involves three transactive processes that ensure effective utilization of limited and distributed knowledge stored in members' individual memory systems. (1) *Updating* the internal skill directory of "who knows what" occurs as members work on interdependent tasks and information about everyone's competencies is observed and stored. This directory updating can be explicit (e.g., conversations) or implicit (e.g., inferred from tasks accomplished). (2) *Allocation* involves routing new, incoming information to the member best positioned to understand and store it. Given the associative nature of human memory, the best member for storing any incoming information is the member who already possesses related knowledge. Such differentiated accumulation of knowledge in local experts leads to *member specialization* (3) *Retrieval* is aided by well-functioning updating and allocation, which minimizes the time required for retrieving knowledge (Wegner, 1995). With multiple retrieval episodes, members come to form shared beliefs about the extent to which they can *credibly rely* on a given member's expertise. These features (specialization and credibility) are used as behavioral indicators of a well-developed TMS (Ren & Argote, 2011).

TMS Model Implementation and Validation

We implement our theory in a multi-agent system that models a six-member collective working on a stream of tasks over 240 simulated timesteps. During each timestep, each individual (or "agent") operates according to the rules we specify to message each other and work on their inbox of tasks. Each task involves resolving a series of problems requiring distinct skills. As agents encounter a task, they decide if they will do it themselves or pass it on to someone else based on the baseline rule-set chosen (see Appendix A2 for all rulesets).

In the TMS baseline, each agent has long-term memory, implemented as a dynamic

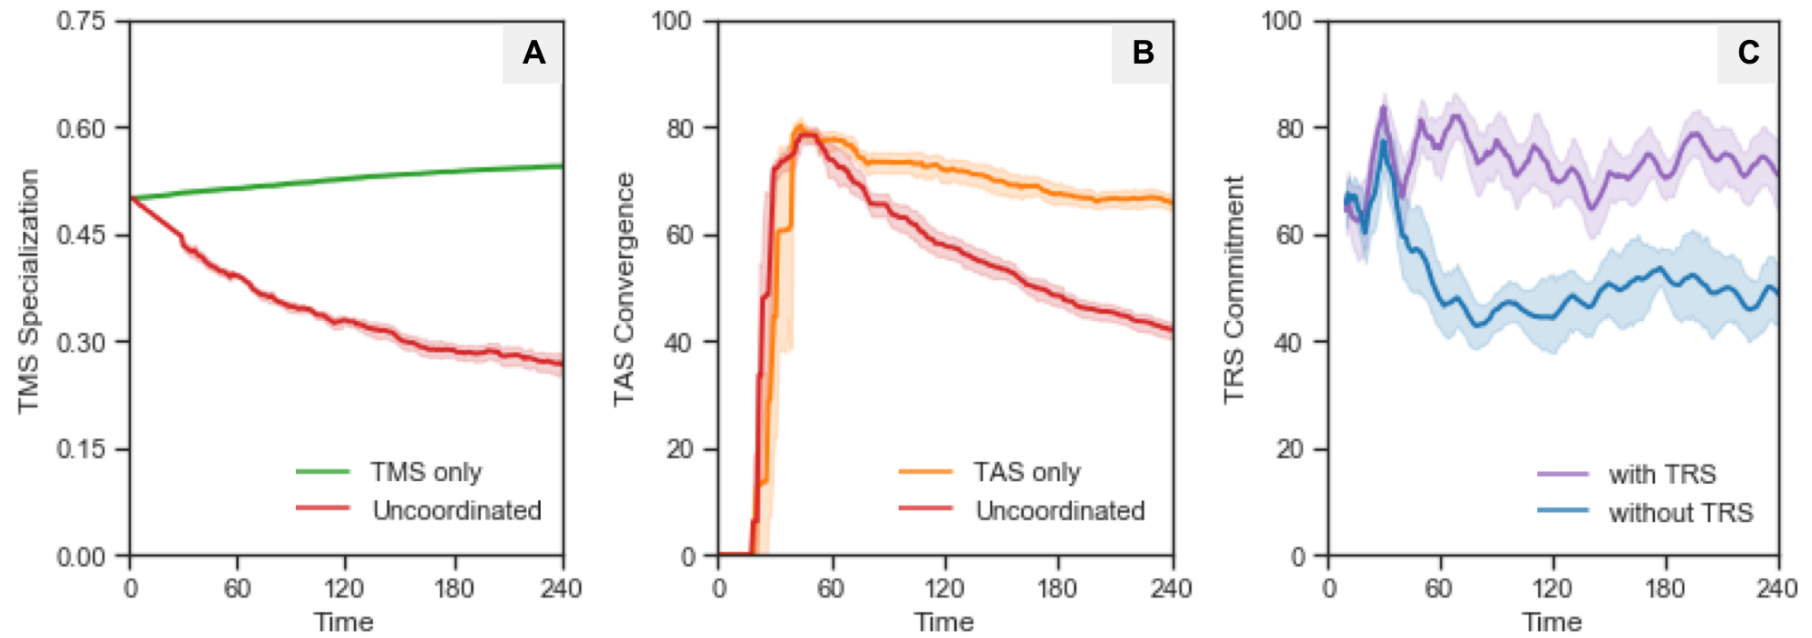
list of skills. Each skill is identified by a label, an expertise level indicated by a number on a scale of 1-10, as well as an indication of the timestep when the skill was last used (e.g. Skill-A::L5::T15 indicates Skill A with an expertise-level of 5 last used at timestep 15). The expertise level determines the amount of work that the member can complete in one timestep. Every time a skill is used for task work, the expertise level increases by a prespecified learning rate (default, 0.05); whenever a skill is not utilized for more than 30 timesteps, “forgetting” occurs which causes the expertise level to decrease by 1. Thus, a member specialized in Skill-A is more productive and improves and maintains their level of expertise by regularly applying their skills. To maintain accuracy of member directories, any increase or decrease in skill level generates an explicit metaknowledge message and is shared with all.

Each agent also has a meta-memory, which is a list of metaknowledge objects, each holding information about a member and one of their skills. Consequently, each agent has a personalized “map” or internal representation of “who knows what” in the collective which manages the allocation, retrieval, and updating processes described previously. The meta-memory map is updated based on messages agents send each other when their expertise increases. This map is used to retrieve information about members and identify who is best to handle a particular task or to assimilate new knowledge on behalf of the group. When running the simulation, we can either allow agents to coordinate using TMS-based information or inhibit it by not allowing meta-memory processes. One simplification we make in the baseline is to assume that metaknowledge is always credible, and focus observations on the extent of specialization (calculated as average within-member skill variance).

On simulating this base case, we observe that when we allow TMS-based coordination, “member specialization” develops as our theory predicts, but is eroded when we block the process specialization (Figure 2.3A). As articulated in TMS theory and demonstrated in the simulation, TMS relies on the logic of skill-based allocation of tasks

Figure 2.3

Model validation for the emergence of (A) TMS, maintains high specialization over time; (B) TAS, maintains high convergence over time; and (C) TRS, maintains high commitment over time.



which delivers value to the collective by efficiently coordinating the storage and use of its distributed knowledge resources. Thus, it is particularly useful when the task involves figuring out and resolving complex knowledge interdependencies. Consistent with this, extant research demonstrates that TMS drives collective performance in a variety of settings by resolving knowledge interdependencies (Lewis, 2004; Ren & Argote, 2011) and is strongly related to psychometric measures of collective intelligence (Kim et al., 2017).

Emergence of Transactive Attention System (TAS)

In addition to memory, the amount of available attention is a primary resource that puts a hard limit on the capacity of any collective to conduct work. Thus, attention management is argued to be central to the role of management in any social system and is at the heart of organizational design decisions (Simon, 1973). Traditionally, organizations have deployed specified workflows, roles, and routines in combination to structure the attention of boundedly rational workers (March et al., 1993). However, more knowledge workers are simultaneously engaged in and juggling their attention across multiple different projects, fluidly transitioning between teams or organizations (Mortensen & Haas, 2018; O’leary et al., 2011; Valentine et al., 2017). This increasing dynamism means that coordination of distributed attention occurs through implicit processes that cannot be managed entirely by traditional structures. To this end, we introduce TAS as a dynamic system comprising of (1) member’s current level of available attentional resources (i.e. individual attention and control system); (2) their current understanding of each others' current focus (i.e. individual meta-attention); and (3) the inter-member processes that facilitate the dynamic updating of meta-attention and subsequent allocation of individual attention and retrieval of joint attention of appropriate members as needed (Figure 2.2).

Individual Attention and Meta-Attention

The individual attention and control system enables individuals to function adaptively

in an information-rich world by selecting the most relevant information to the task at hand while suppressing irrelevant information. Three control processes are responsible for the selection and maintenance of relevant information within the focus of active attention (Knudsen, 2007). The heuristic-driven fast selection process is responsible for automatically filtering the wealth of incoming information based on salient characteristics (Koch & Ullman, 1985). The top-down deliberate selection process is responsible for directing action to increase the flow of relevant information from the external world as well as deliberate recall of relevant information from the individual's long-term stored memory (Egeth & Yantis, 1997). Finally, the maintenance process involves chunking of closely related tasks and information pieces in order to minimize the costs associated with cognitive switching as one transitions between mental actions (Desimone & Duncan, 1995). To complement these processes, individual meta-attention provides globalized cognitive control to help expertly navigate situations involving choice among multiple tasks (Lavie et al., 2004; Miller & Bigi, 1979). Meta-attention is essential for monitoring individuals' workload and guiding their selection of focal tasks without incurring undue switching costs. In this way, meta-attention is critical to the efficient use of attentional resources. We extend this conceptualization of meta-attention to also encode information about others' focal tasks and their priority relative to new tasks that may emerge. Thus a well-developed meta-attention will help coordinate task distribution by allowing members to redistribute workload when it is unbalanced, thereby reducing bottlenecks in the project's workflow. By contrast, poorly-developed meta-attention will result in misallocation of time to tasks and hinder work progress.

Inter-Member TAS Processes

Just as with TMS, TAS involves three inter-member transactive processes that facilitate the efficient utilization of collective attention. (1) *Updating* refers to maintaining a dynamic and accurate understanding of "who is doing what." Effective updating demands

communicating frequently or developing systems (like shared calendars or task trackers) that enable members to indicate their availability or observe each others' focal activities. (2)

Allocation of attention involves routing new tasks to members or redistributing tasks in a manner that minimizes bottlenecks and aligns with collective priorities (which are informed by TRS; see next subsection). Well-functioning attention allocation prevents attention-priority misalignments and results in *member focusing* behavior centered around the most valuable tasks or projects. To maintain this as the number of members in the collective increases, systems for managing attention allocation become increasingly important, as the amount of information to process increases exponentially. In some situations, this might be handled by a single decision-maker (e.g. scrum master) or decentralized and handled asynchronously (e.g. by an issue tracking or shared tagging system). (3) *Retrieval* occurs when joint, synchronous attention is needed from multiple members, such as to discuss complex issues or handle highly interdependent tasks. Poor retrieval management can lead to wasted attentional resources, either by calling for simultaneous attention when it is not necessary and thus distracting from other important work, or when it is needed but weak temporal coordination causes delays and lost time. Akin to the negative costs of cognitive switching associated with individual attention, a well-functioning collective attention system needs to minimize the switching costs incurred or time lost as members working on interdependent tasks attempt to *retrieve* each others' attention to exchange information or coordinate hand-offs. The effective formation of joint attention is foundational to building shared intentionality and working toward a common goal (Tomasello et al., 2005).

Collectives with well-developed attention retrieval processes will exhibit organized patterns of temporal *convergence*, sometimes manifest as "burstiness," characterized by periods with low levels of interaction punctuated by intense periods where members are simultaneously active on shared work and highly responsive to each others' requests. Burstiness in work

patterns has been shown to predict innovation (Riedl & Woolley, 2017) as well as collective intelligence (Mayo & Woolley, 2021).

TAS Model Implementation and Validation

In our agent-based TAS model, each agent has a limited attention span (Kahneman, 1973) consisting of a dynamic list of up to three recently-accessed skills from memory, and is responsible for maintaining the skills associated with the current task in active attention. If the specific skill an agent needs for the current task was not recently accessed then the agent incurs a switching cost by losing one timestep. In this timestep, it locates the skill in memory and brings it to its active attention. Therefore, there is an advantage to maintaining the skills associated with high-priority tasks in active attention.

Each agent also has meta-attention, which is a dynamic list of metaknowledge objects, each holding information about a member and one of their tasks. Consequently, each agent has a personalized internal representation or map of "who is doing what" which helps to manage updating, allocation, and retrieval. The map is dynamically updated by storing information about each member's focal task as tasks move around. The map helps members select the highest priority tasks to focus on, as well as to redistribute tasks among members to ensure coverage of high priority tasks and maximum utilization of available attention. Thus, in a highly functioning TAS, member attention is redistributed to high-priority tasks soon after they enter the system. We can disable TAS-based coordination by inhibiting meta-attention processes. One simplification we make as a baseline is by always providing a clear signal of task priority, and observing the extent members collectively converge attention on higher priority tasks (calculated as the percentage of total completion time spent by members actively working on a task, weighted by task priority).

On simulating this base case, we observe that "convergence" of collective focus on higher priority tasks is maintained at a high level for collectives with TAS, but when we

disable TAS-based coordination this coordinated focus is eroded (Figure 2.3B). As articulated in TAS theory and demonstrated in the simulation, TAS relies on the logic of priority-based allocation of tasks to maximize the value of collective attention. It does so by avoiding bottlenecks (and unused resources) while retrieving and reallocating attention as high-priority tasks emerge or priorities change. Thus, it is particularly critical for collective work under conditions of high workload with continuously evolving dependencies and priorities.

Emergence of Transactive Reasoning System (TRS)

While TMS and TAS work together to ensure efficient deployment of available resources, specifically member skills and focus, they do not guarantee that the collective is selecting the most valuable goals in terms of overall rewards or personal importance to members. To supplement the focus on efficiency and productivity that TMS-and TAS-based coordination can drive, collective reasoning serves an important maintenance function by monitoring the value of the collective's goals, both in terms of relative opportunities in the environment and, relatedly, in terms of members' personal goals and motivations. In the increasingly complex and fluid environment surrounding collective work, where more and more jobs are contract or "gig-based" (Deutschkron & Pearce, 2019; Valentine et al., 2017) and boundaries are highly permeable, collectives need to attract and retain members. A collective that doesn't pursue valuable goals or operates in a manner that ignores members' needs is likely to face higher levels of resource uncertainty, in the form of member disengagement, burnout, or turnover. Thus, tackling task and resource uncertainty is critical to the collective's ability to sustain itself. To this end, we propose TRS as a dynamic system comprising of (1) members' goals and motivational needs (i.e. individual goal system); (2) their current understanding of each others' primary goals and needs as well as the value of alternative opportunities (i.e. individual meta-reasoning); and (3) the inter-member processes that facilitate the dynamic updating of meta-reasoning and subsequently maximizing joint

rewards via dynamic retrieval of member resources and their allocation to collective goals (Figure 2.2).

Individual Reasoning and Meta-Reasoning

The organization and motivational contents of our individual goal system dictate how we juggle our significant and often conflicting goals and how we select the goals we choose to pursue. Goal systems can be commonly understood as a hierarchy of associatively linked cognitive structures that represent a set of goals along with a set of tasks that are their means of attainment. Goals are typically associated with motivational needs that are fulfilled when a goal is achieved. Continuously fulfilling proximal motivational needs is critical for sustained progress towards valued distal goals (Bandura, 1991). While goal hierarchies maximize short-term rewards and illicit goal commitment, individual meta-reasoning is needed to achieve metacognitive control that maximizes longer-term rewards. Meta-reasoning involves monitoring and determining whether to continue, switch strategies, or terminate pursuing the current set of proximal goals to avoid local reward maximas (Ackerman & Thompson, 2017). This input to metacognitive control is critical for sustaining performance when the individual is experiencing psychological uncertainty or the task environment is continually changing. Someone with well-developed meta-reasoning is able to effectively adapt to the changing situation by reconfiguring their goal hierarchies and ensuring reward maximization. In addition to engaging in reasoning about one's own cognitive and motivational states, meta-reasoning facilitates understanding of others' goals and motivations. One way this is achieved is via the development of a theory of mind, wherein one can shift perspectives and infer the cognitive and motivational states of others through observation and interaction (Demetriou, 2000). Theory of mind has come to be understood as a foundation of social intelligence (Woolley et al., 2010) and deficits in this ability are considered to be signs of significant mental disability, such as autism (Baron-Cohen, 2000)

Inter-Member TRS Processes

TRS involves three inter-member transactive processes that facilitate collective reasoning. (1) *Updating* refers to maintaining an accurate view of the goals and motives of others as their local environment changes. Since goals are inextricably linked to the state of their environment, this updating also involves updates to task mental models and facilitates the formation of shared situational awareness (Endsley, 2015). (2) *Allocation* relates to the collective's ability to allocate priorities and achieve goal alignment, with implications for the distribution and allocation of resources to goals. The collective's ability to leverage changes in its task environment is based on the extent to which its members are able to monitor new opportunities and threats as well as its success in integrating this information to realign priorities. Having a diverse pool of highly motivated members is useful for recognizing opportunities (Afuah & Tucci, 2012; A. Carsrud et al., 2009) and shared situational awareness (Saner et al., 2009). The resulting *clarity* about the portfolio of collective goals and their priorities allows for effective guidance of TAS for resource allocation and helps recognize the needed missing resources. (3) *Retrieval* involves the extent to which the collective is successful in garnering the commitment and effort of existing and prospective members (including their time, skills, social capital, etc) in the service of the collective and its objectives. Recent work on transactive goal dynamics has articulated how individuals in interdependent relationships form a single unit wherein they adopt and hold other- and system-oriented goals. When members adopt a dense set of goals that are aligned with the system's goals, then they will be willing to devote more resources towards their joint goals as well as provide goal support (Fitzsimons et al., 2016). This results in higher levels of goal persistence and *commitment* to the collective (Koestner et al., 2012; Moreland et al., 1993). This is supported by research on game theory demonstrating the power of "we-reasoning" mechanisms for persistence on collective goals, even in the face of competing personal goals

(Bacharach, 1999). Thus, a well-developed TRS enables the provision of more resources and clearer goal priorities to TMS and TAS to utilize and execute efficiently.

TRS Model Implementation and Validation

We instantiate a simplified goal reasoning system for each agent in accordance with self-determination theory (Ryan & Deci, 2018). The agent's current goals and subsequent task selection are driven by the extent to which their motivational needs for competence and relatedness exhibit autonomous regulation (high satisfaction) or controlled regulation (low satisfaction) at a given timestep. Unless their individual needs fall below their threshold of controlled regulation, the member autonomously chooses tasks that are aligned with collective goals and exerts maximum effort towards the focal task. If satisfaction falls below their threshold, the member continually chooses tasks that are in line with their personal goals until their satisfaction level rises above the threshold. This is when their goals may fall out of alignment with collective goals. Moreover, the agent exerts only 50% effort thereby producing only half the work warranted by their experience level. In this way, the collective pays a double penalty —misaligned goals and half the effort— if it fails to maintain its members' motivations above their threshold levels. Meta-reasoning is a dynamic list of metaknowledge-objects, each holding information about a member and current goals. Consequently, each agent has a personalized internal representation or map of "who values what" which is dynamically updated by observing other members whenever their contribution level changes. It is utilized to identify members with eroding motivation levels sooner. This gives enough time to provide goal support as well as identify tasks that align with personal goals without detracting from collective goals, and the collective is able to retrieve maximum effort out of them. When running the simulation, we can block TRS-based coordination by not allowing meta-reasoning processes. One simplification we make in this baseline system is always allowing accurate assessment of task priorities. We don't model the emergence of

shared goals via TRS as both, the value of exploration under variable task environments (e.g. Afuah & Tucci, 2012; Hong & Page, 2004; March, 1991) and the processes of the emergence of shared knowledge (Grand et al., 2016) are well studied in the literature.

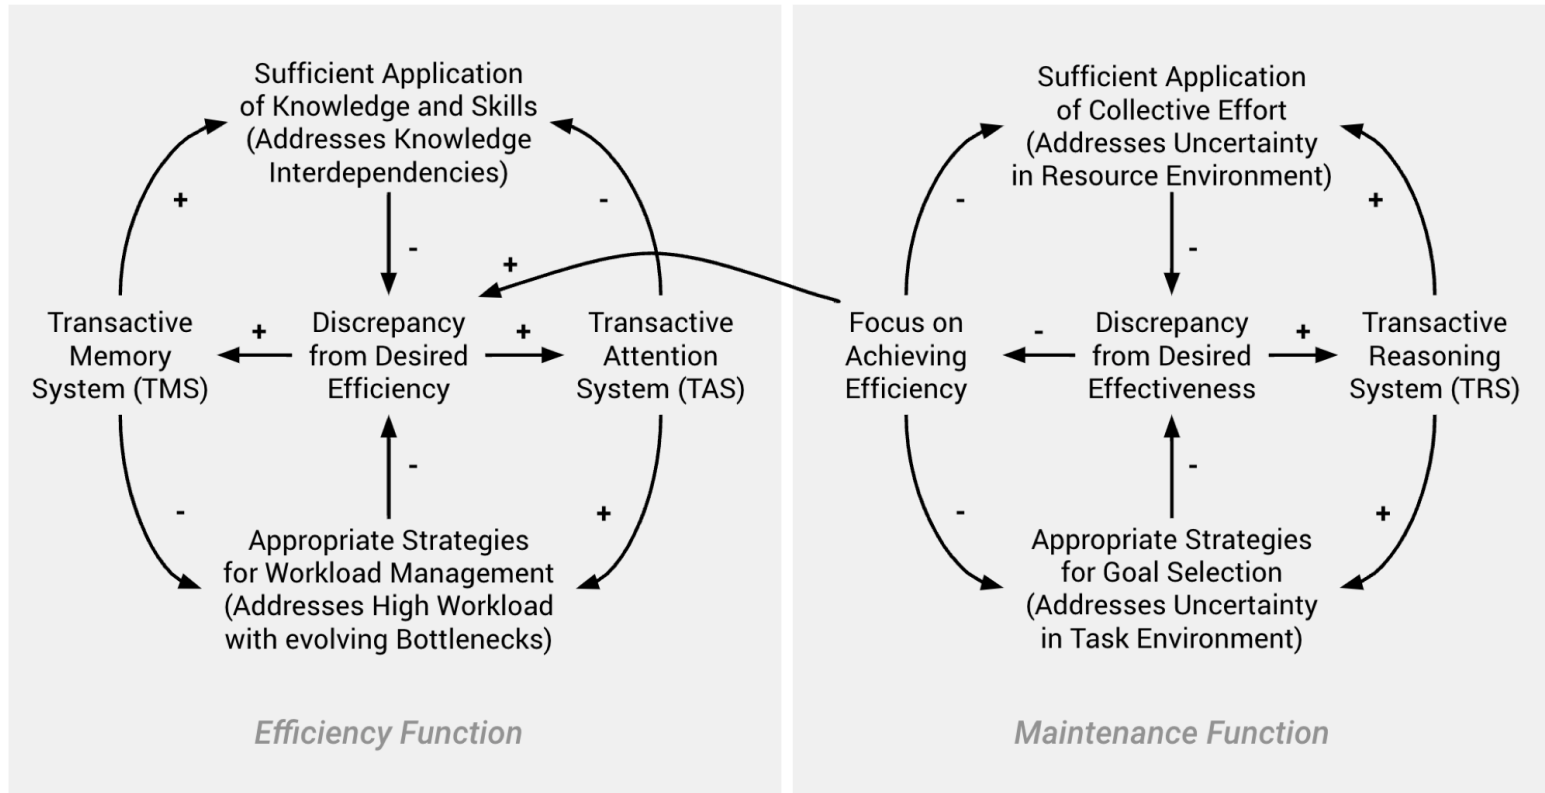
We observe in our simulation that when we allow TRS-based coordination, "member commitment" (10-step rolling average of member effort) is sustained at a higher level, but when we block the process, member commitment erodes (Figure 2.3C). As articulated in TRS theory and demonstrated in the simulation, TRS relies on the logic of reward and resource maximization which delivers value to the collective by aligning individual and collective goals, resulting from the ongoing evaluation of threats and opportunities which shape members' decisions about effort. Thus, a strong TRS is important in environments where the collective is faced with uncertainty in tasks and resources, as the lack of such a capability would result in member withdrawal.

Emergence of Collective Intelligence via Adaptive Regulation of Transactive Systems

In the previous section, we describe the bottom-up processes whereby TMS, TAS, and TRS emerge independently. However, a critical element of intelligence in any system is the ability to adapt to changes in complexity. Adaptation requires that these systems engage in mutual regulation in response to environmental complexity. In this section, we theorize how TMS, TAS, and TRS are locked in multiple feedback loops that trigger each other in response to top-down, contextual factors. Together, they complement each other and dynamically fulfill their efficiency and maintenance function, providing the foundation for collective intelligence (see Figure 2.4 for an overview). For each function, we formally explore the regulatory relationship under varying combinations of environmental threats by extending our agent-based simulation to conduct a virtual experiment. We draw on the results to propose testable hypotheses related to the conditional utility of each system.

Figure 2.4

Overview of adaptive regulation of transactive systems. TMS and TAS regulate each other for efficient utilization of resources (see inside left panel). While TRS regulates the tradeoff between maintaining resources and seeking efficiency (balance resources across panels).



Efficiency Function: Co-Regulation of TMS and TAS for Efficient Resource Use

As described previously, collectives have two essential distributed resources—member skills and attentional focus. Fulfilling the efficiency function is a continuous trade-off between the appropriate utilization of these resources (e.g., the person with the most task-relevant expertise might not be the most available). Because TMS regulates appropriate utilization of knowledge and skills, reliance on a well-developed TMS will be particularly important in environments where there is a high level of knowledge interdependence (Ren & Argote, 2011). However, over-reliance on the logic of skill-based allocation of tasks might also lead to the formation of bottlenecks, where tasks pile up for experts while others' attention is under-utilized. The probability of such bottlenecks goes up as the collective experiences a high workload without a concomitant shift to balancing attentional resources. We posit that in a collectively intelligent, self-regulating system, deficiency in workload management would trigger the collectives' TAS processes and reallocate the most urgent, important, and unattended tasks to maximize the value of attentional resources. At the same time, while a purely TAS-driven logic of coordination can prevent bottlenecks, it may do so at the cost of quality due to inadequate skill matches. To the extent that this causes poor outcomes, a compensatory contribution from the TMS-driven logic of coordination will be triggered. In this way, in conditions of high collective intelligence, TMS and TAS regulate each other until an equilibrium is reached between the quality and efficiency of work given available personnel and the current workload.

Virtual Experiment #1: TMS and TAS Co-Regulation

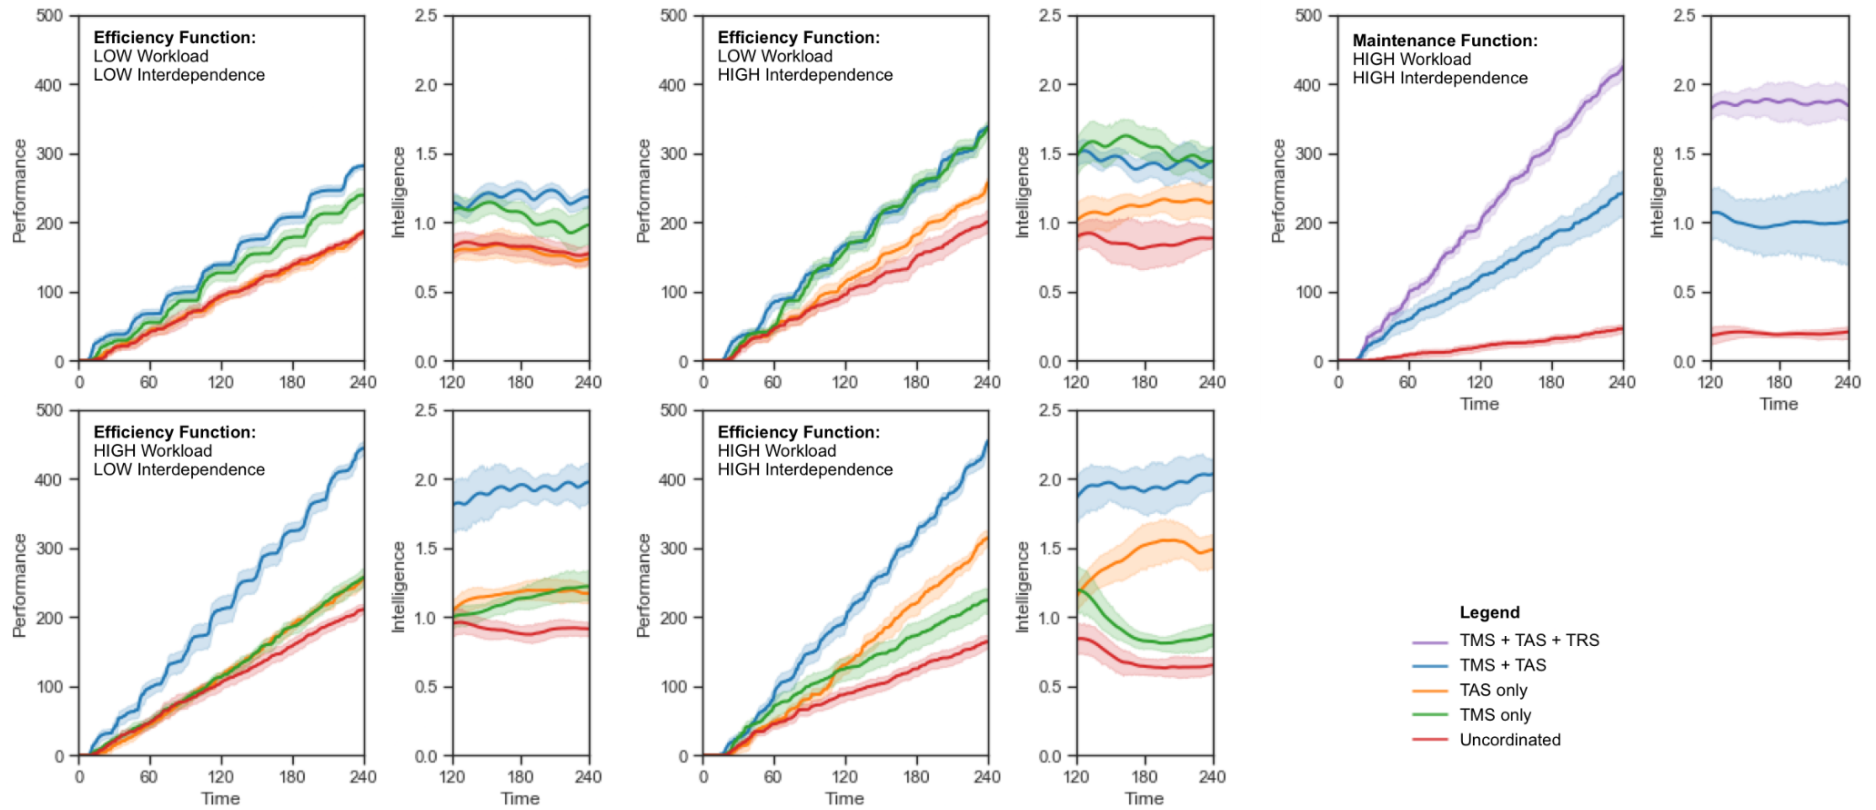
In order to explore the independent and joint effects of TMS and TAS under different combinations of environmental threats, we extend our simulation to test six-member collectives working in four configurations: uncoordinated, TMS-only, TAS-only, TMS + TAS. We simulate the collectives' behavior across two levels of knowledge interdependence

(low vs. high) and two levels of workload (low vs. high). Our virtual experiment was run on each simulated collective 20 times ($N = 4 * 2 * 2 * 20 = 320$ runs). For the purposes of this virtual experiment, we disable agents' individual reasoning systems and TRS; agents have no personal goals and exert maximum effort at each timestep. Knowledge interdependence was manipulated by changing the number of distinct skills needed to complete a task while holding the total amount of taskwork constant. High interdependence trials involved 6/6 distinct skills requiring all members to work on every task before completion. Low interdependence trials required 3/6 distinct skills. In manipulating workload, in low workload trials, we randomly assigned each member a new task every 30 timesteps, while in high workload trials members were assigned new tasks every 20 timesteps. Each task had associated reward points that the collective received on its successful completion; however, the reward was reduced by half for tasks completed after the deadline. The number of reward points associated with a task determined its priority level. Performance was measured based on total points accumulated in a given time period, and collective intelligence was measured as the rate at which points were accumulated; this was calculated as the performance-slope updated dynamically every 120 timesteps.

The results of the virtual experiment demonstrate that collectives with co-regulated *TMS+TAS* outperform *TMS-only* and *TAS-only* collectives across all conditions (see Figure 2.5, Column 1 and 2). We only observed differentiation between *TMS-only* and *TAS-only* collectives in the high interdependence conditions; *TMS-only* collectives performed better in the high interdependence/low workload conditions, while *TAS-only* collectives performed better in the high interdependence/high workload condition (Figure 2.5, Column 2). The most striking evidence for the benefit of co-regulation is observed in the low interdependence /high workload condition, where both the *TMS-only* and *TAS-only* conditions perform similarly, but the co-regulated *TMS+TAS* condition exhibits significantly higher collective intelligence.

Figure 2.5

Virtual Experiment Results: (Top two rows) Efficiency Function: Collectives with co-regulated TMS and TAS outperforms TMS-only and TAS-only Collectives across all conditions. (Bottom row) Maintenance Function: When member effort is at risk, collectives with TMS, TAS, and TRS exhibit higher intelligence than collectives with TMS and TAS only.



In summary, co-regulation of the transactive memory system (TMS) and transactive attention system (TAS) is necessary to determine the most efficient allocation of available member skills and attention. While allocation guided by joint *TMS+TAS* logic is better than either of them individually, the relative reliance on one or the other is conditional on the particular environmental threats faced by the collective. This suggests the following aggregate behavior:

Hypothesis 1. *TMS improves collective intelligence by buffering against the negative effects of knowledge interdependence. Specifically, maintaining a well-developed TMS becomes more important for achieving high collective intelligence in an environment with consistently high levels of knowledge interdependence.*

Hypothesis 2. *TAS improves collective intelligence by buffering against the negative effects of task workload. Specifically, maintaining a well-developed TAS becomes more important for achieving high collective intelligence in an environment with consistently higher levels of task workload.*

Maintenance Function: Regulate TRS for Effectively Exploring & Retaining Resources

Focusing on efficiency in task execution can lead collectives to ignore important changes happening in the broader environment. Important information that could be ignored relates to task and resource uncertainty, such as changes in the environment that alter the value of current vs. alternative goals, as well as changes to the alignment of members' personal and collective goals. Routine disregard of members' goals and needs leads to disengagement, burnout, and eventually turnover (Cappelli, 2020; Zucker, 2019). A collective with a well-developed TRS will monitor the internal and external environment, and trigger action to make changes when needed to increase the value of activities for members individually or the collective as a whole, which would ultimately lead to higher collective intelligence. Navigating this duality between preserving the system state and exploring new

system states (i.e. morphostasis vs morphogenesis) is a classic dilemma of any open system (Buckley, 1967).

Virtual Experiment #2: TRS, TMS, and TAS Co-Regulation

We expand our virtual experiment to explore the additional value that TRS can bring to TMS and TAS co-regulation, particularly in changing or uncertain task and resource environments. In this virtual experiment, we again simulate six-member collectives, this time working in three different configurations: uncoordinated, *TMS + TAS*, or *TMS + TAS + TRS*. Based on the results of the first virtual experiment, we focus on the high workload-high interdependence condition and manipulate resource uncertainty by enabling the individual goal reasoning system which results in varying member effort. Each configuration is run 20 times ($N = 3 * 20 = 60$ runs). As discussed previously in the section on TRS, when a TRS is functioning, individual and collective goals are aligned which leads to higher levels of commitment and effort; when there is no TRS they become misaligned and effort declines, with downstream effects on performance and collective intelligence. Consistent with virtual experiment 1, performance is measured as total points accumulated at a given time while intelligence is measured as the rate at which points were accumulated. As shown in Figure 2.5, results demonstrate that while collectives that solely focus on efficiency (*TMS + TAS*) are better than uncoordinated collectives, they exhibit lower performance and intelligence (see Figure 2.5, Column 3) as compared to collectives that have all three (*TMS + TAS + TRS*). These collectives maintain individual and collective goal alignment, and thus member effort remains high, thereby exhibiting higher levels of collective intelligence. In summary, TRS supplements efficiency-oriented logic by monitoring and maintaining goal alignment and member motivation. This suggests the following effect of aggregated TRS behavior:

Hypothesis 3. *TRS improves collective intelligence by reducing the negative effects of task and resource uncertainty. Specifically, maintaining a well-developed TRS*

becomes more important for achieving high collective intelligence in an environment with higher levels of task and resource uncertainty.

In this way, we theorize that the emergence of collective intelligence can be characterized by the co-regulation of three transactive systems enabling the balance of its maintenance function (which “selects right people and goals”) with its efficiency function (which “gets them done quickly”). The influence of these systems will depend on the key sources of complexity in the environment — specifically knowledge interdependence, workload, and task and resource uncertainty.

Collective Intelligence in Open Source Software Teams:

An Empirical Investigation of Transactive Systems Model

In the previous two sections, we described a transactive systems model of collective intelligence narratively and demonstrated its sufficiency in an agent-based model. We then used virtual experimentation to explore the relative utility of the three transactive systems under critical environmental threats (e.g. well-developed TAS when the workload is high). We now turn to operationalizing and testing our model on empirical data from real-world teams.

Rigorous empirical validation of a complex adaptive system model involves tracking the behavioral dynamics of TMS, TAS, and TRS (forms of growth, decline, etc.) as a function of different environmental threats and comparing them to data from real teams (Schwaninger & Groesser, 2009). Although a simpler test can be conducted by statistically aggregating the behavioral markers of the TMS, TAS, and TRS and environmental threats from real teams’ data and using them to predict the level of collective intelligence exhibited by them. While inadequate for validating causal dynamics, such a statistical test is a reasonable “proof of concept” and adequate for establishing the distinct and contingent role of TMS, TAS, and

TRS under specific environmental threats. Working within the constraints of available behavioral data from real teams, a rigorous test was not possible. Thus, we simplify the model dynamics to test hypotheses that describe the conditional utility of each transactive system in predicting collective intelligence under different environmental threats.

In this section, we analyze 18 months of member activity data from 476 open-source software (OSS) project development teams in the “Node” community. Given the complex, dynamic situations that characterize open-source projects, with high membership turnover and wildly fluctuating levels of member activity, OSS is a prime setting for statistically testing our proposed hypotheses.

Overview of Research Setting: Open Source Software (OSS) Project Teams

OSS projects are a great example of community-based, geographically distributed, virtual work where members contribute to multiple self-organized projects (O’Leary & Cummings, 2007; Vasilescu et al., 2016). OSS projects are typically volunteer-driven; while in some cases programmers are compensated by their employer for time spent on strategically important OSS projects, the vast majority of work is carried out by volunteers who work on any project that interests them at whatever time they choose. The combined effect of members’ geographic and temporal dispersion, differing backgrounds, multi-teaming behaviors, and voluntary membership make coordination uniquely challenging (van Wendel de Joode & Egyedi, 2005). Given the context of OSS projects, membership churn is high, and thus the ability of a collective to retain members by assigning them fulfilling tasks and thus getting them to consistently contribute is particularly critical for maintaining high-quality performance over time. This makes the setting suitable for testing the value of TRS.

Despite the lack of commercial backing, OSS projects cover a staggering range of complex software applications; from operating systems (e.g. Linux) to browsers (e.g. Firefox). Entire programming ecosystems (e.g. Python, Node.js) including the framework and

several thousands of smaller libraries or packages used within the framework are developed and maintained as OSS projects. A vast majority of OSS projects are hosted on GitHub, a code-management software. Contributors to an OSS project primarily contribute by writing software code to add new features or improve existing features of the project's software. They often fix bugs that have been reported by engaging in comment-based conversations. GitHub provides an infrastructure for storing and organizing all of the different pieces of a project, in addition to a system for reporting issues and managing taskwork via "tags". This makes available for analysis a host of archival project activity data like contributors, code commits, reported issues, comments, as well as micro-level task management events like labeling, assignments, referencing, etc. for a given code repository. Thus, the ability to observe and analyze a project's characteristics as well as the contributors' time-stamped digital activity traces through their publicly logged events provides an exciting opportunity to test a variety of questions, including our transactive systems model.

OSS contains a variety of projects with varying sizes, complexity, and development activity belonging to a variety of ecosystems, each with idiosyncratic community norms, making it a good setting to test the impact of TMS and TAS. Controlling for variation in community norms is particularly important for our analysis, as we operationalize our transactive systems based on emergent behaviors. Too much variability within the sample can make it a nuisance variable. Therefore, we restrict our archival analysis to a single ecosystem: "Node.js". Unlike many OSS communities, the official node website periodically reports on quality, maintenance, and popularity scores for each project as a way for its users to choose between projects that serve similar functions. This makes it a good setting to evaluate a project's success in maintaining high-quality code.

Data and Operationalization of Constructs

For this analysis, we focus on 476 projects from the Node.js ecosystem on GitHub.

The 476 projects in the final sample are filtered down from about 400K possible node projects with GitHub data (as of July 2019). The filter ensured that each project in our sample was actively developed by a team (at least 3 contributors per month between Jan 2018 and July 2019, 18-month analysis window) as well as actively used by the community (at least 1000 downloads in the preceding year, signaling maintenance of high-quality code). Next, we created measures of environmental complexity and level of transactive systems use by aggregating theorized behavioral markers observed in the 18-month member-activity data.

Environmental Complexity

We first considered the critical sources of environmental complexity that threaten a projects' ability to maintain high-quality code — knowledge interdependence, workload, and task and resource uncertainty. In selecting OSS as a context for testing our model, we determined that task and resource uncertainty was a factor that was uniformly high across all projects, given that contributors are free to enter and exit at will and the members are free to negotiate and choose project goals. Thus there was no means of or need to operationalize a metric to capture the level of task and resource uncertainty; nonetheless, a TRS to manage this source of environmental complexity would be critical for all projects.

Workload. The number of issues to be handled varied considerably by project, even among projects of a similar size. As mentioned, all tasks including issue reports were logged into an issue tracker, making it quite straightforward to capture how much activity was taking place in the project. *Workload* was operationalized as the total number of open issues per member on the project within our analysis window.

Knowledge Interdependence. The structure of these software projects was highly emergent and complex. As different contributors initiate a feature, they create code that they can choose to add to an existing file or create a new file. Resolving an issue by updating the code often necessitates understanding, changing, and reorganizing code across multiple files.

On one hand, as the number of code changes increases the level of knowledge interdependence among the co-edited files also increases. On the other hand, the code-base tends to get organized into relatively independent modules that reduce knowledge interdependence. To calculate this, we first identified the number of code modules in each project by constructing a file co-edit network and counted the number of clusters of closely connected files that tend to get changed together. Next, we calculated *Knowledge Interdependence* based on the average number of edits (or commits) per module (commit count / code modules count). Thus, larger values indicate projects where a typical module has undergone a relatively higher number of edits creating stronger interdependencies with them. Failure to understand how focal code change affects code in other interdependent files can lead to adding potential bugs that break the project, reducing its quality and making it hard to maintain.

Transactive Systems and Indicative Patterns

We draw on the indicative patterns theorized for each of the transactive systems to operationalize indicators of each in the sample (see Figure 2.2).

Transactive Memory System (TMS). Two indicators of the level of TMS in a collective are whether or not members have credibility in their areas of expertise, and whether they specialize by concentrating their work in those areas. While we did not have access to participants' perceptions to gauge credibility, data on work patterns were available to provide insight on specialization. In this context, specialization reflects the extent to which members focus their contributions on fewer vs. across more code modules. To evaluate this, we first constructed a member co-edit network for the project and, using cluster analysis, identified different clusters or contributor groups within the project. Then, we created an index of *TMS-Specialization* by comparing the number of contributor groups and the number of code modules (contributor groups count / code modules count). If each module is worked

on by a unique contributor group then $TMS=1$ (i.e. highly specialized), else $TMS<1$ indicates the extent to which groups contribute across code modules (i.e. less specialized).

Transactive Attention System (TAS). Two patterns that are indicative of TAS function are the degree to which members are focusing on tasks that are collective priorities, and convergence around temporal patterns of work by responding to one another and work to resolve issues in quick succession. The issue tracker which serves as the main hub for OSS projects facilitates discussion via comment threads and also enables issue tagging. Tags are used to direct other members' attention by indicating the priority level of an issue or alerting specific members to provide input. As members communicate to resolve issues and contribute code over time, Github passively logs the project's activity data. Using these data, we calculated two different indicators of TAS. *TAS-Focusing* is operationalized based on the amount of tagging activity per member. Because convergence occurs when activity takes place in more bursty patterns (i.e. activity concentrated in shorter time periods vs. distributed more broadly across time), *TAS-Convergence* is operationalized as the extent of temporal burstiness (calculated as fano factor) of the activity stream.

Transactive Reasoning System (TRS). Two patterns that are indicators of the level of TRS functioning are collective goal clarity and member commitment. We did not have access to data from participants which would enable us to estimate goal clarity. However, our data does capture the extent to which a contributing member persists in working on and continuing to contribute to the project. To operationalize this, we created a member-by-month contribution matrix such that each cell was marked 1 if a member contributed during that month. Then, we created another persistence-matrix by taking a 3-month rolling average across months for each contributor and multiplying it (element-wise) with the contribution matrix. Each cell in this matrix denotes each member's persistence index, ranging from 0 (no contribution in the past three months) to 1 (contributed in each of the past three months).

TRS-Commitment is the average of the persistence index across all contributors across all months. Higher values indicate the project's ability to retain their contributing members.

Collective Intelligence (NPM-Score)

The node.js community analyzes the activity of all the projects on a rolling basis and reports scores across three dimensions (quality, maintenance, and popularity) by aggregating different aspects of the project over the year. For example, maintenance denotes if the project is being actively maintained and depends on metrics like release and commit frequency, proportion, and duration of open issues over the past year. Similarly, quality is reported by combining more complex software features like code coverage of automatic tests, the extent of outdated dependencies, and the release of stable versions (Lichtman, 2018). *Collective intelligence (NPM Score)* is constructed by averaging the 'quality' and 'maintenance' scores reported by the community. These scores are driven by the core team's ability to keep the project code operational, up-to-date, and bug-free over time in the face of evolving membership, workload, and code complexity.

Control Variables

We create four control variables: *project-capacity* calculated as the average count of contributors each month; *project-age*, months since first code commit (first member contribution); *project-size*, operationalized as the total number of commits; and *project-popularity*, indexed based on scores reported by the community.

Results

Descriptive statistics and correlations for all study variables are displayed in Table 2.1. First, inspecting the zero-order correlations among study variables reveals some important observations relating to sources of environmental complexity in this setting. While we note that *workload* demonstrates an expected and significant negative relationship with collective intelligence ($r = -.31; p < .01$), *knowledge interdependence* is not correlated with

collective intelligence ($r = -.05$ $p > .01$).

To test our conditional hypotheses, we standardized all variables and ran linear regression models in two steps. In the first model, we regress on the *environmental complexity* variables to determine which of them are a dominant threat (significant negative effect) that hinders the OSS teams' ability to maintain high-quality code over time, i.e. exhibit collective intelligence. Then, we add *transactive systems* variables to examine their conditional utility in buffering against these dominant threats.

Hypothesis 1 predicted that TMS improves collective intelligence by reducing the negative effects of knowledge interdependence. Knowledge interdependence does not appear to be a significant threat to maintaining the quality of OSS projects (Table 2.2-Model 1). Accordingly, we see that TMS is also not a significant predictor of collective intelligence (Table 2.2-Model 2). This doesn't mean that OSS teams do not need TMS. Recall, we theorize that TMS and TAS co-regulate each other, and together, give rise to efficient use of resources. In fact, we observe that OSS teams in our sample exhibit well-developed TMS (mean = 0.8; sd = 0.3; theoretical range = [0,1]) but cannot expect it to significantly predict collective intelligence unless the corresponding environment threat it buffers against is negative and significant. Our inability to test the dynamics is a limitation of the metrics that capture aggregated behavior over our time window.

Hypothesis 2 predicted that TAS improves collective intelligence by reducing the negative effects of task workload. As demonstrated in Table 2.2-Model 2, and in support of hypothesis 2, both indicators of TAS (focusing and convergence) moderate the effect of workload on CI. In decomposing the interaction, we see in Figure 2.6 that the effect of TAS was particularly dramatic for teams experiencing very high levels of workload. Finally, hypothesis 3 predicted that TRS improves collective intelligence by reducing the negative effects of task and resource uncertainty. As mentioned previously, task and resource

Table 2.1. *OSS-GitHub Study Descriptive Statistics and Correlations*

	Mean	SD	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]	[9]	[10]
[1] CI (NPM Score)	0.86	0.11	-									
[2] Workload	17.99	15.23	-0.31	-								
[3] Interdependence	29.91	29.17	-0.05	0.12	-							
[4] TRS Commitment	0.64	0.11	0.23	-0.33	0.08	-						
[5] TAS Focusing	109.85	102.80	0.13	0.24	0.23	0.38	-					
[6] TAS Convergence	62.17	73.48	0.06	0.26	0.18	0.08	0.47	-				
[7] TMS Specialization	0.80	0.30	0.00	0.08	0.02	0.16	0.20	0.18	-			
[8] Project Popularity	0.32	0.20	0.04	0.36	-0.03	-0.44	-0.03	0.16	0.06	-		
[9] Project Size	1844.8	2799.5	-0.05	0.28	0.58	0.14	0.40	0.39	0.27	0.08	-	
[10] Project Capacity	15.43	17.85	0.00	0.34	0.20	-0.11	0.30	0.36	0.23	0.34	0.53	-
[11] Project Age	54.94	24.58	-0.20	0.24	0.12	-0.18	-0.05	0.02	0.14	0.32	0.23	-0.04

Note. N = 476; $p < .01$ in bold; CI = Collective Intelligence

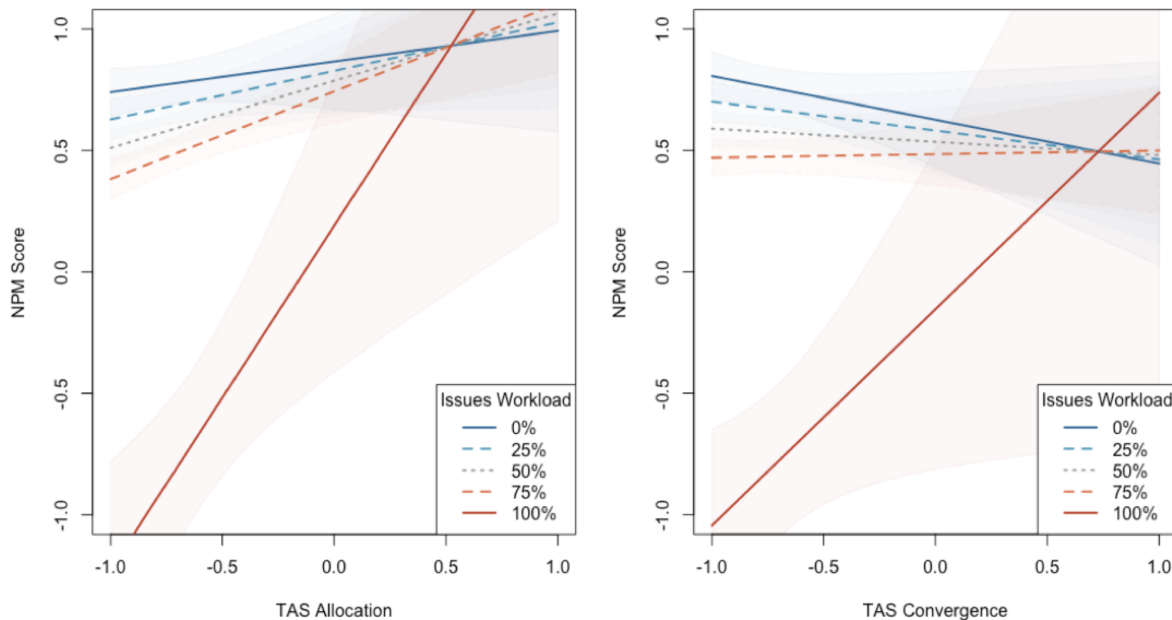
Table 2.2. *Test of Hypotheses in OSS-GitHub: OLS Regression Results*

	Collective Intelligence (NPM Score)			
	Model 1		Model 2	
	Coeff	SE	Coeff	SE
Environmental Complexity				
Workload	-0.613 *	-0.081	0.130	-0.193
Interdependence	-0.063	-0.085	0.000	-0.086
Transactive Systems				
TRS Commitment			0.135 *	-0.057
TAS Focusing			0.694 *	-0.195
TAS Focusing * Workload			0.606 *	-0.260
TAS Convergence			0.440 *	-0.178
TAS Convergence * Workload			0.607 *	-0.256
TMS Specialization			0.025	-0.072
TMS Specialization * Interdependence			0.034	-0.086
Controls				
Project Popularity	0.204 *	-0.045	0.288 *	-0.047
Project Size	0.170	-0.093	-0.069	-0.102
Project Capacity	-0.024	-0.078	-0.020	-0.077
Project Age	-0.191 *	-0.044	-0.132 *	-0.043
Intercept	0.153 *	-0.076	0.768 *	-0.144
N	476		476	
R ²	0.161		0.245	
adj R ²	0.150		0.224	

Note. * $p < .05$

Figure 2.6

High workload projects without a well-developed TAS suffer the most. Collective intelligence is operationalized as NPM score - an index capturing the projects' ability to maintain good quality code.



uncertainty in the form of loss of contributors is uniformly high in the OSS context, and thus not an element that varied across projects. As shown in Table 2.2, TRS, as indexed by member commitment, has a significant positive relationship with collective intelligence. The findings support hypothesis 3. In the OSS context, where member contributions are at risk, having a well-developed TRS positively predicts collective intelligence.

General Discussion

Although the view of organizations is shifting from well-structured machine-like entities designed for performance to dynamically responsive complex adaptive systems capable of exhibiting collective intelligence, we have lacked a detailed process model describing how collective intelligence emerges and is sustained under different conditions.

The purpose of this paper was to articulate and evaluate a minimally viable socio-cognitive architecture—Transactive Systems Model—of collective intelligence. Building from classic theory in management (March et al., 1993) we posit that, in the absence of rigid structures, boundedly rational actors can leverage their metacognitive capabilities to coordinate their limited and distributed memory and attention as well as their diverse goals to overcome their cognitive limits as a collective. Building on this, we described the bottom-up processes that lead to the emergence of transactive memory, attention, and reasoning systems (TMS, TAS, and TRS) and articulated the persistent behavioral patterns that are indicative of their development.

We instantiated these processes in an agent-based simulation to conduct an initial sufficiency validation along with two virtual experiments to explore how the three transactive systems mutually regulate and adapt based on the critical sources of environmental complexity, thereby exhibiting collective intelligence. Results from the first experiment demonstrated that collectives utilize their resources—members’ skills and focus—most efficiently when TMS and TAS co-regulate each other. Specifically, we hypothesized that the relative reliance on one or the other is conditional on the level of threat associated with different environmental complexities faced by the collective. Using the second virtual experiment we hypothesized that, when resource uncertainty is the most threatening source of environmental complexity, member commitment is at risk. In these conditions, high collective intelligence results in the deployment of TRS to dynamically ensure the maintenance of collective resources rather than allowing the pursuit of efficiency to dominate. Finally, we empirically tested these hypotheses by analyzing archival data from 476 open-source software teams. Data and observation demonstrated that this is a context where resource uncertainty (i.e. member commitment) and workload are the largest threats, And consistent with predictions, evidence supported the conditional value of TRS and TAS.

By contrast, knowledge interdependence was not a source of complexity in this context, and as predicted we did not observe an influential role for TMS.

Implications

This work has important implications for both theory and research. In terms of theory, our process-oriented framework and computational model provides a precise conceptualization of the emergent and regulatory dynamics (Figure 2.2 & 2.4) that shape how and when collective intelligence develops. By pairing narrative theory with agent-based modeling, our paper exhibits the benefits of theories of emergence prevalent in team cognition (Carley, 2001; Kozlowski & Klein, 2000). In addition, our complex adaptive systems approach enables a more holistic perspective where we can investigate the co-regulatory effects of our theorized transactive systems. This has important implications for research on social cognition where the literature has accumulated evidence for independent effects of key constructs (DeChurch & Mesmer-Magnus, 2010) but lacks studies that test their joint or co-regulatory effects. In our virtual experiments and our empirical data, we find evidence of such co-regulation.

Next, we extend the literature on social cognition by consolidating existing work to articulate two new constructs, TAS and TRS. While the attention-based view (ABV, Ocasio, 2011) of the firm takes inspiration from individual attentional processes and articulates the selective, executive, and vigilant aspects of attending to information flowing from the environment, TAS builds upon them by leveraging the concept of meta-attention to articulate how it helps members coordinate each others' focus. In this sense, ABV functionally overlaps the processes involved in team situational awareness and the subsequent emergence of shared goals (Burke et al., 2006; Endsley, 2015; Rico et al., 2019). Additionally, complementing extant work on transactive goal dynamics (Fitzsimons et al., 2016), TRS consolidates relevant constructs to articulate the emergence of a system's maintenance function. While our

initial model of TRS dynamics proposed in the paper is adequate for our research questions, more explicit integration of these constructs would be useful in future studies.

In terms of implications for research, in addition to articulating our transactive systems theory of collective intelligence and providing initial validation, we also identified operational metrics for evaluating TMS, TAS, and TRS in empirical data. This facilitated the study of these constructs in an empirical setting to replicate and extend our model. In addition, the indicative patterns associated with each transactive system provide the basis for additional computational process metrics to enable the study of these constructs in other settings (Riedl et al., 2021).

Finally, our work has important implications for research in information systems. The transactive systems model can be used to further the basic goals of collective intelligence research—to detect, diagnose, and enhance collective functioning. Identifying the contingencies associated with sources of environmental complexity is important for targeting critical areas of focus. The indicative patterns associated with each of the transactive systems provide a starting point for diagnosis, and could likely be augmented by technology-based systems. Such insights could, in turn, provide starting points for interventions. However, given the highly interdependent nature of complex adaptive systems, it is important to think of the system holistically to avoid unintended consequences where an intervention in one system creates a new problem in another (e.g. Gupta & Woolley, 2018).

Limitations And Future Directions

Parsimony is a central tenet of theorizing complex adaptive systems (Carley, 2001; Miller & Page, 2007). As a result, our model makes various simplifying assumptions which may limit its applicability to some settings and are fertile grounds for future extensions. First, while metacognition is central to our theorizing we have not built in allowance for deficiencies in this area; as modeled, our agents have an accurate awareness of self and

others' mental states. However, we know that individuals vary significantly in their abilities to pick up on others' mental states, which has significant implications for collective intelligence (Engel et al., 2014; Meslec et al., 2016; Riedl et al., 2021). Furthermore, as the communication cost for achieving "complete metaknowledge" will increase exponentially with the size of the collective, both inaccuracy and variability of metacognition across members is likely to further increase. Therefore, future developments in this work will need to test the boundaries of the accuracy and variability of metacognition and consider what structures might be necessary such as subgroups/clusters or hierarchy to enable the critical level to be achieved.

Second, our model ignores emotions. Emotion and cognition are deeply intertwined (Dolan, 2002). Emotion plays a central role in reasoning, not only in decision making (Schwarz, 2000) but also in collective reasoning by affecting the extent to which members care about each other's personal goals and are willing to trust and coordinate with each other (Druskat & Wolff, 2001). Thus, inter-member emotional dynamics can have a profound impact on transactive processes regulating collective attention and memory, especially under high-stress situations. In addition, we chose to focus on sources of environmental complexity that have been widely demonstrated as threats to coordination. However, in further developing the model, additional considerations such as sense-making under varying levels of ambiguity and risk preferences (March & Olsen, 1979; March & Shapira, 1992; Weick, 1995) would add important nuance and improve generalizability.

Finally, the simplified "proof of concept" test of the conditional utility of different transactive systems in the context of OSS teams collapses the theorized dynamics and simply examines the relationships by aggregating behavioral markers of all the transactive systems as well as collective intelligence. Such regression tests of emergent phenomena suffer from multiple weaknesses. Of particular importance are two limitations: (a) its inability to

disentangle the direction of causality between the threat and the corresponding transitive system, and (b) its susceptibility to washing out effects of a threat that may have been high for only a small proportion of the time window it was aggregated over. Therefore, future work that maps the theorized dynamics with temporal observations that span across various levels and combinations of environmental threats would be better suited for testing this systems theory.

Conclusion

This paper introduces a transactive systems model for the emergence of collective intelligence and presents a multi-method systems approach for investigating its adaptive behavior. Our theory articulates the metacognitive and inter-member processes underlying the emergence of collective memory, attention, and reasoning, and suggests how these transactive systems co-regulate in response to different task constraints to sustain a high level of collective intelligence. In so doing, this work theorizes a role for metacognition in the place of traditional organizational structures, where members are aware of their limitations and leverage the collective to overcome them. Much like the cognitive architectures that have guided the development of artificial intelligence (ACT-R, Anderson & Lebiere, 1998; SOAR, Laird, 2012; Newell, 1990), this socio-cognitive architecture holds the potential to be further formalized in computational terms to explore the dynamics of collective intelligence and pinpoint opportunities for enhancing it. Such guidance could prove very timely —advances in robotics and AI are unlocking new and exciting organizational forms that leverage their individual strengths. In order to unleash the true potential of human-machine collaboration, we need to design with an understanding of the emergent and regulatory dynamics necessary to cultivate collective intelligence.

III

Dynamic Indicators of Collective Intelligence:

The Diagnostic Role of Collaborative Processes in Transactive Systems Model

The Transactive Systems Model of Collective Intelligence theorizes that collective intelligence results from the emergence of transactive attention (TAS), memory (TMS), and reasoning (TRS) systems that are a result of individual-level cognitive processes and member interactions regulating one another in response to changing environmental conditions. While computational models or retrospective archival analyses can provide the comprehensive information necessary to analyze the development of the theorized systems and their consequences (such as the analysis in Chapter 2), in many settings it would be helpful to have observable indicators, available in real-time, to signal the level of functioning of different parts of the system. Such indicators would enable managers or even automated systems to identify problems as they develop along with potential points of intervention.

In order for such indicators to have practical utility, they would need to be valid signals of the regulatory functioning of the underlying transactive cognitive systems—TAS, TMS, and/or TRS. In addition, they would need to be both more specific and functionally “upstream” from the collective intelligence indices for the entire system. The observation of a decline in some overall metric of collective intelligence is important, but typically incorporates too many inputs at too high a level and too late in the process to be able to pinpoint a specific problem with the team. Thus, useful indicators would be more specific and available early enough to course correct difficulties in a particular part of the system before the problems are amplified across the system as a whole. In complex adaptive systems terms, this indicator acts as the ‘sensory mechanism’. It perceives changes in the environment as well as reflects how efficient and well-maintained the system is, at the current time.

Here, we build on ideas originally introduced into the traditional teams literature to identify essential collaborative processes which can serve as dynamic indicators of the level of collective intelligence developing in a team or other collective. We first formalize our argument for the value of a well-developed sensory mechanism by briefly discussing four problems that any collective experiences when continually adapting to a changing environment. Next, we describe three collaborative processes based on models in the extant literature and theorize how they are related to various transactive system processes, thereby claiming their value as a decomposable sensory mechanism. In the next section, we build on an agent-based model presented in prior work (see chapter 2) and present a virtual experiment to explore how our three theorized collaborative processes develop as signals of the level of regulatory functioning of the three transactive systems. In a further extension of our model, we also develop a composite indicator of collective intelligence using these collaborative processes, which serves as a useful real-time proxy for performance-based measures of collective intelligence that are typically available at a later stage.

In developing these diagnostic indicators, this work lays the foundation for translating the socio-cognitive architecture of collective intelligence underlying our descriptive Transactive Systems Model into machine-readable, dynamic indicators that can be used for real-time diagnosis and intervention—by human managers, or by AI systems.

Four Common Problems That Plague the Management of Complex Adaptive Systems

For theorizing a socio-cognitive architecture of collective intelligence (Transactive Systems Model), we relied on two important features that govern the dynamic behaviors of complex adaptive systems: the *emergence* of transactive systems from localized member interactions and *adaptive regulation* of transactive systems in response to a changing environment. Similarly, informing our ability to design and manage complex adaptive

collectives requires understanding two more features— *non-linearity* and *path dependence*— that characterize these dynamic behaviors. Both of these issues point to the critical role played by indicators in successfully diagnosing and intervening in complex adaptive systems. In this section, we identify and briefly describe four common problems posed by non-linear and path-dependent behaviors.

First, a system where heterogeneous members respond to their local conditions tends to find itself locked in multiple reinforcing and balancing feedback loops (Figure 2.4). Viewed from the global perspective, adaptive behavior arises as the relative strengths of these feedback loops shift. These shifts can be prompted by external perturbations as well as the specific nature of how the feedback loops are connected internally. The shifting dominance of feedback loops gives rise to abrupt or *non-linear* changes in behavior. Two common forms of non-linear changes include a small perturbation resulting in a disproportionately large output and a larger perturbation having little to no effect on the system output. Thus, collectives are prone to exhibiting plateauing as well as tipping point behaviors (Meadows & Wright, 2008; Miller & Page, 2007). Longer gaps between action and feedback in a non-linear system may be a recipe for undesirable outcomes. A collective with an *insufficient sampling of changes* will fail at sensing maladaptive behaviors and often find itself experiencing unintended consequences. A particularly egregious form of non-linearity arises when a balancing feedback loop has a delay built into it. It makes the system output susceptible to cyclical and sometimes hard to detect oscillations (Arrow et al., 2000; Meadows & Wright, 2008). Delays can trick managers into generating *ill-calibrated responses to changes* in their environment and get collectives stuck in a cycle of reactionary and overcompensatory actions that get worse as time passes. The right types of diagnostic indicators can enable better sensing of problems and effective responses to counteract the tendencies to either overcompensate or wait too long to respond.

Second, most complex systems go through some degree of non-linear change—reinforcing loops, compounding returns, lock-ins— which results in a corresponding level of path dependence (Sterman and Witt, 1999). *Path dependence* refers to the ways in which early activities and decisions open up or close off possibilities for future action. A system's history uniquely constrains the probability of successfully pursuing specific actions in the near future (Miller & Page, 2007; Page, 2006). A particularly important aspect of a system's history that has received special attention is its 'input or initial conditions' (Meadows & Wright, 2008). A collective whose members have all of the skills needed to address the tasks at hand as well as have well-defined roles and norms for developing transactive systems appropriate for the workload can quickly achieve high collective intelligence. On the other hand, many collectives are likely to start out with some form of *compositional or structural handicap*, lacking appropriate skills or poorly thought-out norms. Such collectives need a way to quickly sense and overcome their situation or risk low collective intelligence behavior. Furthermore, the early developmental trajectory of transactive systems also creates path dependence. For example, experiencing early success with TMS-based task allocation can lead a collective to encode heuristics and routines centered on reliance on expertise. If this occurs, it can become harder for a collective to successfully sense and adapt to changes that require a different logic of coordination; in other words, a collective can become locked-into TMS-driven *regulatory inertia*. This may lead the collective to have difficulty adapting to higher workloads and consequent bottlenecks and fail to notice that for some tasks, relevant expertise is less important and they should use a TAS-based resource-driven regulatory logic. In this way, the level of intelligence exhibited by a collective is highly sensitive to the initial conditions and rules of interaction that shape their path, particularly in the absence of diagnostic indicators available at earlier stages to mitigate the tendency toward inertia.

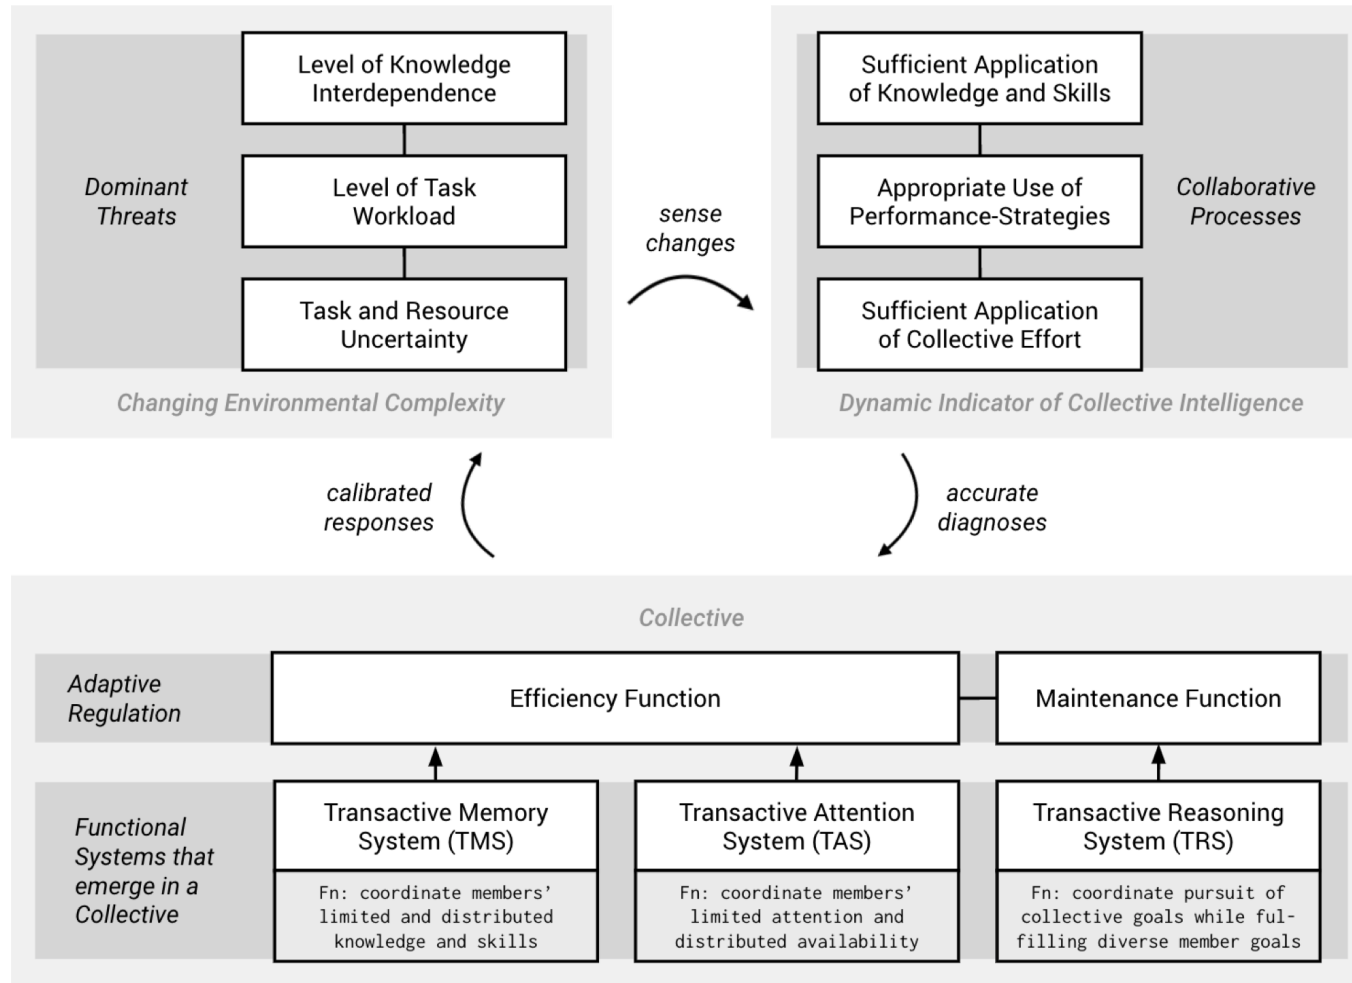
In addition to valid and early detection of these four types of problems, effectively dealing with non-linearity and path dependence requires diagnostic indicators that isolate a particular part of the process or the system as a likely focus for intervention. And, because a complex adaptive system is continuously changing, it is important that the indicators also update dynamically to ensure an accurate assessment of an issue as well as the impact of an attempted intervention. In the next section, we describe three collaborative processes we have identified from extant literature and their relationship to the transactive systems included in our socio-cognitive architecture of collective intelligence. We argue that these relationships make the three collaborative processes ideal diagnostic indicators for collectives in dynamic settings (see Figure 3.1 for an overview).

Collaborative Processes and their Relationship with Transactive Systems Model

In extant work on collective intelligence, typically indicators are generated based on performance on multiple tasks, perhaps in different environments, that vary in complexity. However, such measures often are only possibly relatively late in the process, leaving little opportunity for adjustment or improvement. Herein lies the core challenge of designing and managing teams that aspire to operate at high levels of collective intelligence. While punctuated use of after-action reviews, retrospectives, and performance feedback has tremendous value, what can we do about course-correcting in real-time, when the work is being done? Using diagnostic indicators can not only be useful for real-time interventions but also facilitate evidence-based performance evaluations. The team performance literature has a fairly established history of research on team processes which have been looked to in various settings as proximal indicators of group functioning. A number of classic models have identified different processes; one widely recognized and empirically supported model is the Normative Model of Team Effectiveness (Hackman, 1987; Hackman, 2011) which argued

Figure 3.1

Collaborative processes dynamically sense the changing environmental complexity (indicator of current collective intelligence) and help the collective accurately diagnose the problem and allow managers to take actions that desirably regulate transactive systems (calibrated response).



that three collaborative processes serve as essential and predictive indicators of team effectiveness. These include: (1) the *sufficiency of the effort* members are expending in carrying out collective work; (2) the *appropriateness of performance strategies*, that is, ongoing choices that members make about what and how the collective work is getting done; and (3) the *sufficiency of the knowledge and skill* the members are applying in carrying out collective work. Subsequent work has validated the use of self-report measures of these processes in longer-term teams at critical transition points (Wageman et al., 2005), as the focus of interventions administered by skilled team coaches (Hackman & Wageman 2005) and, more recently, as a differentiator of teams that are high versus low in collective intelligence (Riedl et al., 2021).

Recognizing the value of initial conditions when managing complex adaptive systems (Hackman, 2012), as well as the need for adaptation to maintain performance over time (i.e. high intelligence) reinforces the value of using dynamic measures of collaborative processes that can serve as diagnostic indicators. Without such inputs, even the best initial conditions can morph into sources of regulatory inertia and poorly sense changes in their goals, norms, and membership evoking ill-calibrated responses. Based on extant research, the collaborative processes— collective effort, performance strategy, and use of knowledge and skill— to identify problems and opportunities in real-time is likely to be of tremendous value for ensuring high levels of collective intelligence. In addition to identifying what indicators might have useful signal value, it is important to determine how they relate to the underlying transactive memory, attention, and reasoning systems in order to determine if and how to make a change to a system, thereby fulfilling the role of ‘sensory function’ in a complex adaptive system (Figure 3.1). As described in the Transactive Systems Model, the three systems engage in mutual regulation in response to environmental complexity, via a network of reinforcing and balancing feedback loops that trigger each other continually until

efficiency and maintenance tradeoffs are reached. We propose that the three collaborative processes are indicative of the ongoing efficiency and maintenance tradeoffs attained by the system at any given time (refer to Figure 2.4 for the regulatory model). These three collaborative processes don't operate in isolation and monitoring them in combination provides real-time information about the functioning of the underlying, largely unobservable complex adaptive system.

Proposition 1. *Collectives with high intelligence are better at monitoring their collaborative processes— collective effort, performance strategy, and use of knowledge and skill— for sensing undesirable changes and diagnosing the associated transactive systems that are causing the problem.*

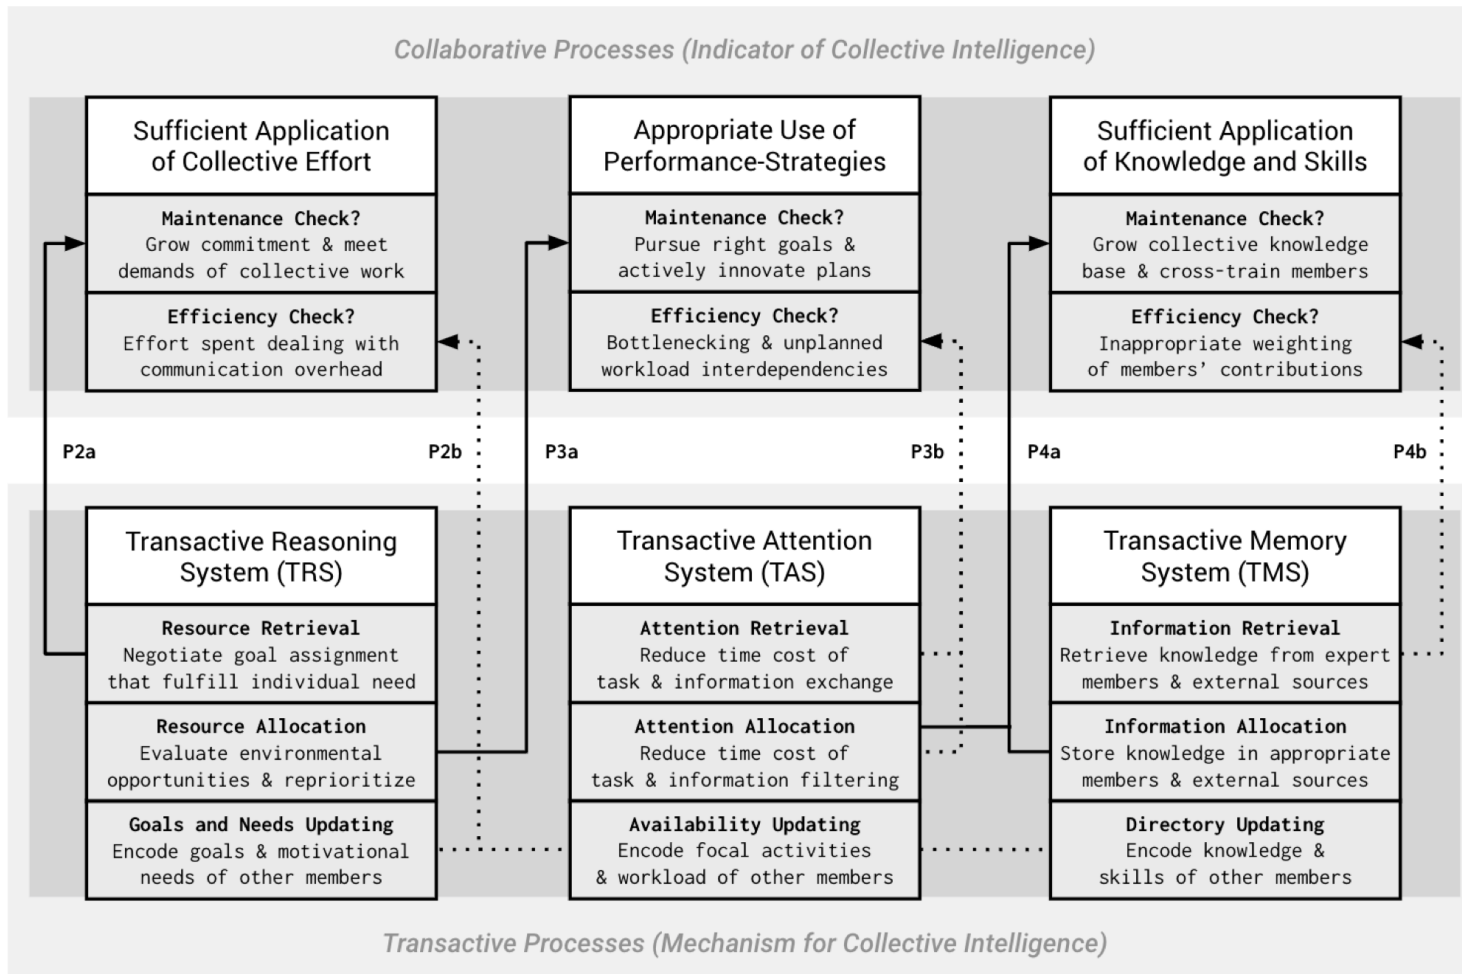
Here we further describe each collaborative process and its relationship to the transactive processes that determine the regulatory tradeoffs necessary for fulfilling the system's efficiency and maintenance functions (see Figure 3.2 for an overview). This framework guides how changes in a diagnostic indicator can be linked to an underlying set of issues associated with the functioning of different transactive systems. Consequently, these indicators can help identify 'leverage points' (Meadows & Wright, 2008) for improving collective intelligence via targeted intervention. In the next section, we will extend the agent-based simulation model from the previous chapter to explore the theorized diagnostic role of these processes. In addition, we also combine the collaborative process metrics to develop a composite metric and demonstrate its value as a leading indicator of collective intelligence by comparing it to the observed trajectories of collective intelligence.

Sufficient Application of Collective Effort

Collaborators are most likely to work hard to accomplish collective goals if they are engaged in work that is meaningful to them, consequential to the collective, and sufficiently

Figure 3.2

Collaborative processes can be used to triangulate and diagnose possible problems with different transactive processes, the underlying mechanism for collective intelligence. Taken together, the collaborative processes indicate the current level of collective intelligence.



challenging while also giving them a substantial level of operational autonomy (Hackman & Oldham, 1976). The resource retrieval process in a well-functioning transactive reasoning system (TRS) involves inter-member negotiations such that the members adopt a dense set of goals that are not only aligned with collective goals but are also motivationally engaging for the individuals. The resulting level of commitment and goal persistence contribute to increasing the total effort members are willing to put towards the collective. The alignment of goals and associated commitment also reduces the feeling of diffused responsibility thereby preventing the most common ailment plaguing work in larger collectives —social loafing. Thus, a drop in collective effort suggests a poorly functioning TRS retrieval process. Digging into which may reveal a lack of focus on fulfilling member needs due to overemphasis on TMS and TAS or poor TRS updating process.

That said, members cannot be infinitely motivated to increase effort and devote all their time to the focal collective's work. In addition to work-life balance, modern workers are often juggling multiple projects across multiple teams and organizations (O'Leary et al, 2012). It would be reasonable to mutually agree upon an aspired level of effort expected from each member (e.g., hours per week). Comparing the combined capacity with a benchmark for sufficient capacity is necessary for accurately diagnosing if the collective needs to focus on motivating members or adding more members. In the absence of individual aspiration levels, the collective can also use effort data from its past behavior as the baseline. Understanding these maintenance trade-offs is essential as member turnover is the biggest source of uncertainty in the resource environment. Moreover, frequently changing or adding members comes with its own complications, a non-linear increase in the communication overhead.

Proposition 2a. *A consistent decrease or constant fluctuation in collective effort suggests poorly developed TRS retrieval processes (maintenance check)*

The most egregious form of efficiency loss experienced by collectives is the effort spent dealing with communication overheads (e.g. poorly run meetings). It directly cuts into the time and energy devoted to productive work. In the transactive systems model, members spend time communicating about each others' skills, workloads, plans, and goals (all three metacognitive updating processes). This understanding forms the necessary basis for the emergence and development of the three transactive systems. As such this is unavoidable and members have to find ways to efficiently maintain the accuracy of metaknowledge to reduce wasted time and effort due to poor TMS, TAS, and TRS updating.

Proposition 2b. *A stable but low level of collective effort devoted to task work suggests poorly developed TMS, TAS, and TRS updating processes (efficiency check).*

Appropriate Use of Performance-Strategies

“Performance-Strategy” refers to two distinct sets of collective choices made by members about (1) which goals are worth pursuing and (2) how will we go about performing the tasks involved in the pursuit of these goals. Successfully navigating the former is essential for navigating high levels of uncertainty in the task environment where rewards associated with existing goals can quickly erode and new opportunities often arise. The resource allocation process in a well-functioning transactive reasoning system (TRS) involves the collective's ability to allocate priorities and achieve goal alignment. It monitors for significant shifts in environmental threats and integrates this information to realign priorities. This allows the collectives that develop truly original or insightful ways of formulating and selecting better goals while leveraging their diverse pool of highly motivated members. The resulting clarity about the portfolio of collective goals and their priorities allows for effective guidance of TAS. Thus, inappropriate use of performance strategies suggests a poorly functioning TRS allocation process. Examining this in combination with collective effort may reveal that the collaborators are not spending enough time exploring and evaluating new

opportunities and are no longer in agreement about goals and priorities due to a myopic focus on TMS- and TAS-driven efficiency.

Proposition 3a. *Consistent use of inappropriate performance strategies suggests poorly developed TRS allocation processes (maintenance check).*

Once the goals are chosen and the plans for pursuing them are made, deploying an appropriate strategy involves figuring out member interaction that results in little "slippage" in the execution of these performance plans. Both the attention allocation and retrieval processes that are responsible for the emergence of a well-functioning transactive attention system (TAS) contribute to effective management of the collective's workload. It involves routing new tasks to members or redistributing tasks in a manner that minimizes bottlenecks and aligns with collective priorities. It also ensures swift resolution of unplanned task interdependencies by facilitating synchronous discussion of complex issues, thereby reducing inefficiencies involved in information exchange. Thus, a diagnostic indicator capturing a decline in the execution of performance strategy points towards a poorly functioning TAS allocation and retrieval process. Examining this in combination with an indicator for use of knowledge and skill (discussed next) can alternatively suggest if the collective is overfocusing on expertise at the expense of efficiency, revealing an over-reliance on TMS or an inability to detect bottlenecks or hidden workload interdependencies.

Proposition 3b. *A decline in the execution of performance strategies suggests poorly developed TAS allocation and retrieval processes (efficiency check).*

Sufficient Application of Knowledge and Skill

A collective is most likely to bring sufficient expertise to bear on its collective work when its composition provides the right number and mix of members who have the knowledge and skills needed for achieving the chosen collective goals. Collectives operating with frequently changing goals often come across a lot of new information and regularly need

the development of new types of skills and the resolution of complex knowledge interdependencies. The information allocation process in a well-functioning transactive memory system (TMS) involves routing new, incoming information to the member best positioned to understand and store it. This fosters the development of local experts who possess the unique ability to perceive and solve highly specialized problems or might develop differential experience in tackling certain niches of an ongoing project. Moreover, the pattern of interactions involved in the information allocation process between members with differing expertise further fosters cross-training. Thus, a drop in a diagnostic indicator capturing the growth of task-relevant knowledge and skills points toward a poorly functioning TMS allocation process. Digging into which may reveal poor TMS updating processes or an inability of any member to interpret and learn the new information that is unlike anything the members have handled before. In that case, it might be meaningful to add another member to the team. Understanding these maintenance trade-offs is essential as adding a new member, especially one that doesn't speak the same language comes with its own complications as previously discussed.

Proposition 4a. *A consistent decline in the level of collective knowledge and skill suggests a poorly developed TMS allocation (maintenance check)*

Finally, one of the most common forms of coordination error committed by collectives is inappropriately weighting members' input. Members are prone to cognitive biases that skew the heuristics they use to assign credibility to others' contributions. Often, attributes related to demographics, personality, or organizational status override the value of subject matter expertise. The information retrieval process in a well-functioning transactive memory system (TMS) involves the identification of subject-matter experts and minimizes the time required for accurately retrieving required knowledge. With multiple retrieval episodes, members come to form shared beliefs about the extent to which they can credibly

rely on a given member irrespective of their other attributes. Thus, a drop in the indicator for appropriate application of knowledge and skill points toward a poorly functioning TMS retrieval process. Triangulating this with diagnostic indicators for appropriate use of performance strategy may reveal a deadline-chasing pattern of behavior that is indicative of over-reliance on TAS or an inability to detect hidden knowledge interdependencies.

Proposition 4b. *Inappropriate application of knowledge and skill suggests a poorly developed TMS retrieval (efficiency check).*

Identify Leverage Points by Putting All of Them Together

In practical terms, diagnostic indicators are unlikely to spell out obvious solutions in isolation. In fact, it is unlikely any collective would achieve uniformly high levels across all three collaborative process indicators, even under an ideal scenario, as the underlying systems are locked in a mesh of reinforcing and balancing feedback loops creating non-linear tradeoffs. The three processes don't operate in isolation. This also means that we cannot directly or independently nudge a specific collaborative process and expect uniform improvement in collective intelligence. A recent study suggests that attempts to do this yield mixed results (Gupta et al., under review). Thus, the most fruitful course of action is to identify all the indicator processes that drop over time, generate, and triangulate possible candidates for intervention by diagnosing the corresponding transactive process using the relationships theorized. This would generate a shortlist of intervention candidates. But, how do we rank them?

Given the regulatory relationships between transactive systems, maximizing the probability that a team exhibit high collective intelligence necessitates achieving the most rewarding efficiency and maintenance trade-off. Hence, we propose that the most effective regulatory scenario would involve focusing on maximizing the leading indicator that combines all three collaborative processes. Then one can rely on the components of this

composite indicator to identify and rank the shortlisted intervention candidates to maximize the possibility of pushing the system from its current level of collective intelligence to a higher level. Such an approach honors the messiness of complex adaptive systems exhibiting non-linearity and path dependence where outcomes are being driven by a large number of interacting variables. It offers us theory-driven diagnostic guides for designing and managing collectives.

Proposition 5. *Interventions that maximize the joint value of all three collaborative processes indicators will be better than interventions that exclusively focus on maximizing any single indicator.*

Exploring Diagnostic Use of Collaborative Indicators via Simulation

We build upon the agent-based simulation that instantiates the Transactive Systems Model (presented in the previous chapter) by operationalizing and monitoring the three collaborative indicators. We then conduct a virtual experiment to (a) observe how the composite indicator derived from collaboration indicators compares with the ideal measure of that capture rate of performance and (b) observe how collaborative indicators behave in the absence of specific transactive systems as well as provide a basis for diagnosis by drawing on theorized relationships. Agent-based modeling (ABM) excels in exploring emergent, non-linear, and path-dependent behavior of complex adaptive systems. ABM is a powerful theory-generating mechanism, rather than serving as a vehicle for empirical testing of theorized relationships (Krackhardt, 2000). We wish to explore the co-regulation of these systems and the correspondence of diagnostic indicators by systematically switching off one transactive system at a time, as well as validate our intuition that these indicators don't operate in isolation. We want to explore how switching of one transactive system triggers

other processes involved in the regulatory loops compounding the problems experienced by the collective over time. Thus, underscoring the complexity of diagnosis and intervention.

Model overview and Virtual Experiment Design

We implement the transactive systems theory in a multi-agent system that models a six-member collective working on a stream of tasks over 240 simulated timesteps. During each timestep, each individual (or “agent”) operates according to the rules we specify to message each other and work on their inbox of tasks. Each task involves resolving a series of problems requiring distinct skills. As agents encounter a task, they decide if they will do it themselves or pass it on to someone else based on the baseline rule-set chosen (see Appendix A2 for all rulesets).

We design our virtual experiment to explore how the three collaborative process indicators behave in the absence of specific transactive systems. For this we simulate four different configurations of our six-member collectives: (1) *TMS + TAS + TRS (all)*; (2) *TMS + TAS (no TRS)*; (3) *TMS + TRS (no TAS)*; and, (4) *TAS + TRS (no TMS)*. To examine the criticality of all three transactive systems for achieving collective intelligence, we subject these teams to the most complex environmental conditions. That is, for all simulation runs, tasks require 6/6 distinct skills, requiring all members to work on every task before completion (high knowledge interdependence), members were assigned new tasks every 20 timesteps (high workload), and members regulate their effort based on satisfaction of their individual needs (high resource uncertainty). Each configuration is run 20 times ($N = 4 * 20 = 80$ runs). The emergent and adaptive behavior of the three transactive systems was validated previously (presented in the previous chapter). No new agent interactions were added. For this virtual experiment, the goal is to operationalize and examine the collaborative process indicators under different conditions to enable us to make comparisons and draw diagnostic inferences.

Indicators Operationalization

Level of Collective Effort. We aggregate the number of work-units each member generates at every timestep to operationalize effort. This is based on their current level of need satisfaction. When fully satisfied, each of the six members is expected to devote 100 work-units per timestep (600 units total). The benchmark for sufficiency of the collective effort is set at 500 work-units at every timestep. This leaves us with some slack. The dynamic level of collective effort is calculated by taking a moving average of total work-units generated during a given timestep and dividing it by 500 (range [0-1]). As the communication is assumed to be instantaneous in the model, there is no effort wasted in terms of coordination overhead. The only two ways to see a drop in effort are if the members were idle or they were not working very quickly due to lack of motivation.

Efficiency of Performance Strategy. The external task environment is held constant across configurations and iterations. That is the overall workload (randomized set of tasks) is the same and each team has an equal opportunity to generate rewards by completing them on time. However, since all teams are subjected to a high workload, there is value in prioritizing more valuable high reward tasks over low reward tasks. Under high workload, the most efficient strategy for achieving higher performance is to make sure that the high reward tasks are not stuck in bottlenecks. We operationalize bottlenecking by calculating the standard deviation across the number of tasks in each member's inbox at a given timestep. Bottlenecking is high if one member is assigned the majority of the tasks, while it is low if the workload is evenly distributed. The measure is created by calculating $\exp(-(bottlenecking / k)^d)$ to get a range between [0-1] where 1 is no bottleneck and the measure tends to 0 as bottlenecking increases. k is a constant used to specify the rate at which the measure drops and can be chosen in benchmarking against acceptable levels of bottlenecking.

Use of Knowledge and Skills. At every timestep that a member is actively working on a task using a specific skill, we calculate the discrepancy between the member's current skill level and the level of skill required by the task. If the members' skill level is equal to or higher than the task requirement, the measure would have a value of 1. In all other cases, it is 0 indicating poor use of a members' knowledge and skills. Finally, the number of distinct skills needed for task work is held constant during the run (6 skills) so there is no opportunity to learn new skills.

Composite Process-based Indicator of Collective Intelligence. The leading indicator of collective intelligence derived from the dynamic values of the three collaborative process indicators is created by multiplying the three measures. As all of them range between [0,1] with 1 representing a high score on each, so consequently the composite metric falls within the same range.

Performance-based Measure of Collective Intelligence. Each task had associated reward points that the collective received on its successful completion; however, the reward was reduced by half for tasks completed after the deadline. Performance was measured based on total points accumulated in a given time period, and the performance-based collective intelligence was measured as the rate at which points were accumulated; this was calculated as the performance-slope updated dynamically every 120 timesteps.

In order to observe the aggregate value and overall trends of the collaborative process indicators as well as compare their composite with the performance-based measure of collective intelligence, we plot the 120-period moving average for all the indicators.

Examining Results and Exploring the Diagnostic Role

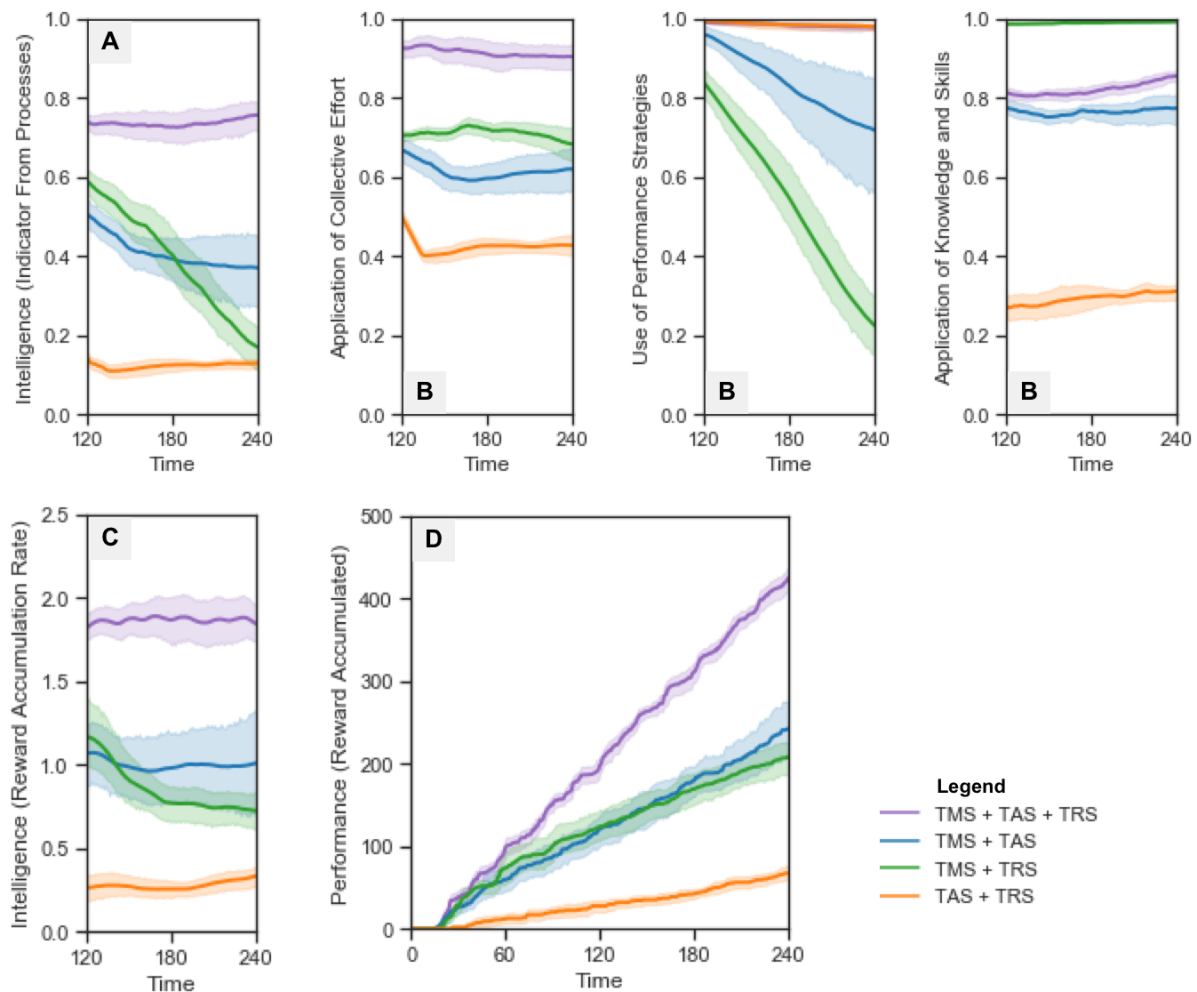
As could be reasonably expected, the condition in which all three transactive systems were allowed to emerge (*TMS + TAS + TRS*) developed the highest level of collective intelligence and best aggregate performance (see the purple line in Figure 3.3, Panel C and

D). We observe that the pattern of results across the four conditions in the performance-based measure of collective intelligence (Figure 3.3, Panel C) is very similar to the composite indicator derived from the three collaborative processes (Figure 3.3, Panel A). This supports our theory-based claims and allows us to use the composite metric as a leading indicator of collective intelligence in real-time as well as leverage its component measures to diagnose the regulatory state of different parts of the system and triangulate the transactive processes that are good candidates for intervention.

Next, we analyze the patterns in the collaborative process indicators exhibited by the three remaining configurations in the simulation, where one specific transactive system is not contributing. Because all of the metrics range from 0-1, we add a reference line at 0.75 to guide inference; indicators at or below this level, or trending downward warrants our attention for triangulating the underlying problem. First, the *TMS + TAS (no TRS)* configuration shows a stable but lower level of collective intelligence as compared to *TMS + TAS + TRS* (see blue lines in Figure 3.3). We observe that the teams in this configuration show a stable but low level of *collective effort* (0.6). Without the ability to develop a TRS this outcome matches our expectation that goals are not aligned or member needs are not being met. The potential for goals to fall out of alignment is constrained as the simulation doesn't model uncertainty in the task environment. Thus, the consequence of a dysfunctional TRS is only restricted to member dissatisfaction. The *application of knowledge and skill* indicator is at par with the best-case scenario while we are seeing a slow erosion of the indicator denoting *appropriate use of performance strategy*. Taken together, we can diagnose that the blue teams have likely prioritized efficiency over satisfying their members. And the resulting drop in the collective effort is just shy of the minimum effort needed to keep up with the high workload. As a result, we are seeing a slow but consistent increase in bottlenecking. If this continues, we will likely see a nonlinear drop in performance in the near future.

Figure 3.3

Virtual Experiment Results: Composite, process-based indicator of collective intelligence (A) along with three underlying collaborative process indicators (B). The Performance-based measure of collective intelligence (C) along with underlying performance score over time (D).



Next, the *TMS + TRS (no TAS)* configuration shows a lower level of collective intelligence compared to *TMS + TAS + TRS*, which is further declining at an alarming rate (see green lines in Figure 3.3). This trend is clearly evidenced in the collaborative process indicator denoting the *use of performance strategy*. Without the ability to develop a TAS this pattern matches our theory-based expectation that members are good at coordinating work in

accordance with collective priorities. The increasing bottleneck has no impact on the *application of knowledge and skill* which continues to be maintained at the highest level. On the other hand, *collective effort* is stable but has declined below the reference line. Taken together, we can surmise that the green teams are overemphasizing TMS-based coordination at the expense of efficiency. Moreover, the existence of bottlenecks makes it easier for members to engage in social loafing such that they prefer to remain idle rather than working on tasks that are important for the collective but difficult for them to accomplish. Left unchecked we would likely observe a further drop in their *collective effort*. With some delay, the lack of work is also going to demotivate members thereby decreasing the amount of effort they are willing to exert when called upon. Such a team would benefit from figuring out “Which tacit norms are blocking members from working on high priority tasks?” and “Who is the bottleneck? What can they do to help them?”

Finally, the *TAS + TRS (no TMS)* configuration shows even more dramatically reduced levels of collective intelligence compared to *TMS + TAS + TRS* (see orange lines in Figure 3.3). This trend is clearly evident in the *application of knowledge and skill* and *collective effort* indicators. Without the ability to develop a TMS the decline of *application of knowledge and skill* matches our expectations but the stark decrease in member effort is surprising. The lack of bottleneck denoted by *appropriate use of strategy* indicator suggests that the orange teams are focused on spreading the workload at the expense of finding the member with the right expertise to do the work. Taken together, we can diagnose that the poor skill-match observed throughout the team has created a situation in which members are taking longer to complete a task. As members are not allowed to abandon tasks in the middle (a feature unique to this simulation), longer work periods on poorly-matched tasks have quickly eroded their motivation. Once demotivated, these members down-regulate the amount of effort they devote to these tasks which further exacerbates (non-linearity) the

speed at which work gets done. This particular combination of TAS, TRS, and simulation rules created a dominant negatively reinforcing feedback loop. Such situations are extremely hard to reverse without significant intervention. Teams, where members are allowed to abandon projects in the middle, are likely to not show such extreme motivational burnout. Such a team will benefit from abandoning their current project in favor of tasks that are aligned with their area of expertise. Particularly, move to tasks that give them a quick win to restore their sense of competency and accomplishment.

General Discussion

Complex adaptive systems, by their very nature, are extremely difficult to observe and accurately diagnose for the purposes of intervening to improve sustained levels of performance. In this paper, we build on the socio-cognitive architecture of collective intelligence and theorize three observable collaborative processes— level of collective effort, appropriateness of performance strategy, and use of knowledge and skill— to be reflective of the ongoing regulatory state of the transactive systems. By theorizing their contribution to efficiency and maintenance functions, I propose that these observable processes can serve as diagnostic indicators and provide real-time information about the *non-linear* and *path-dependent* behaviors of the underlying, largely unobservable complex adaptive system. We explored these collaborative process indicators in an agent-based simulation and investigated their general utility in signaling the level of regulation achieved by different transactive systems. We also examine their validity as a real-time leading indicator of collective intelligence by combining the individual process indicators and comparing the resulting metric to the summative performance-based assessment of collective intelligence. The real-time indicator offers the benefit of enabling monitoring and diagnosis of underlying problems along with pointers to the transactive systems potentially involved in difficulties.

Implication for Leadership and Coaching Functions

The results highlight that the collaborative indicators rarely change in isolation. Even if only one transactive system is initially the source of dysfunction, the regulatory relationship over time creates secondary effects that erode the overall collective intelligence of teams as a result of multiple factors. Left to their own devices, most complex adaptive systems tend towards dysfunction. Cultivating high collective intelligence requires constant vigilance and maintenance, often performed by a leader who can serve as a “coach” who proactively intervenes to address issues. This is consistent with extant work on team coaching, and the importance of this function for providing effective leadership (Hackman & Wageman, 2009):

“Teams that expend sufficient effort on the work, deploy performance strategies that are well-aligned with task requirements, and bring ample knowledge and skill to bear on the work are quite likely to perform well. By contrast, teams that operate in ways that compromise their standing on these three processes are likely to underutilize their collective resources and turn in suboptimal performances. Competent coaching helps members work together in ways that foster process gains and minimize process losses for each of these three performance processes.” (p.18)

There are two major hurdles to be overcome if the leadership or members of a collective intend to succeed in cultivating high intelligence. The first is developing a heightened awareness and deep understanding of how the member interactions determine their ability to perform. And second is developing competence in identifying ‘leverage points’ in real-time, and creatively designing well-targeted, small-scale interventions or nudges that can transform even a large social system. Great coaches, consultants, or leaders may be naturally gifted to do so. We posit that the Transactive Systems Model, together with its relationship to the collaborative indicators theorized in this paper, can help explicitly define a procedure for monitoring, diagnosing, and intervening in the functioning of complex adaptive social systems (see Figure 3.1 for an overview of the full model). This not only reduces the barrier to entry but also opens up interesting possibilities for partly or fully

automating this process. Next, we discuss the possible functional roles an artificially intelligent virtual coach (in the form of a dashboard, automation, or algorithm) can play in improving a team's collective intelligence.

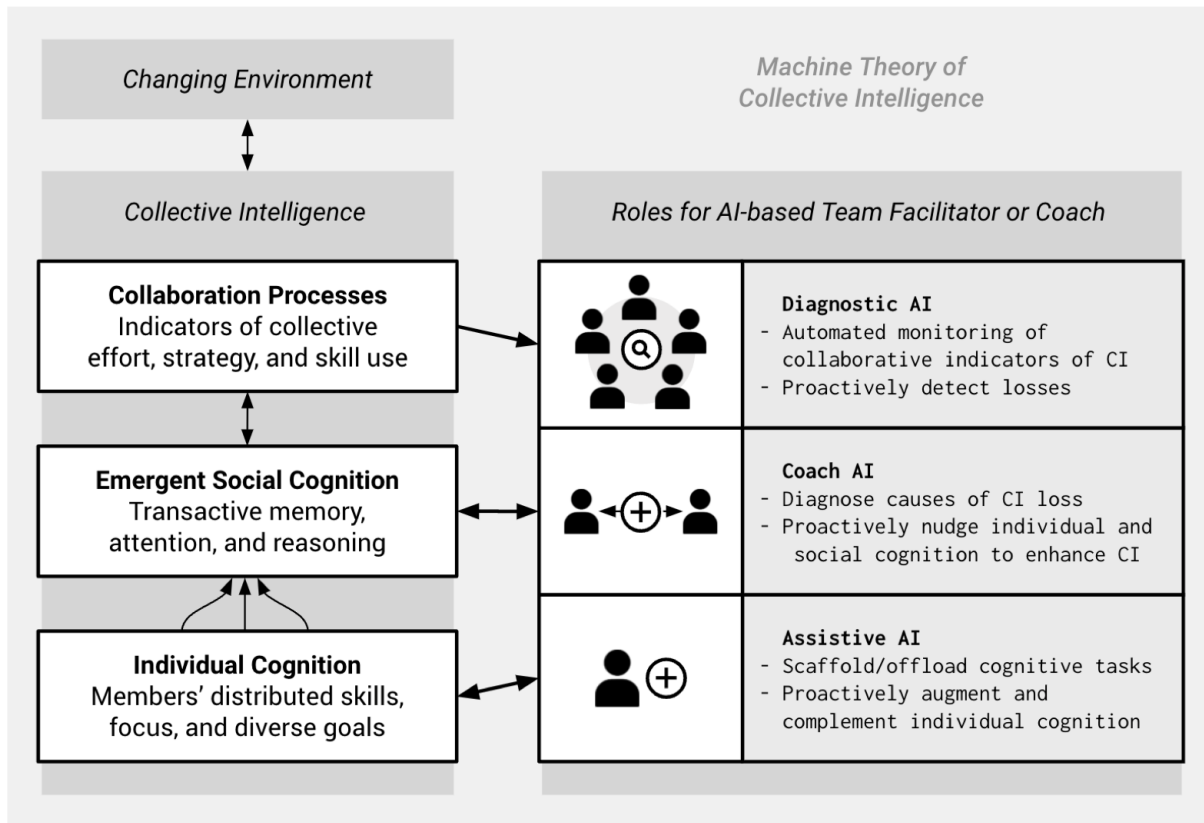
Foundations for the Development of AI-based Team Facilitator or Coach

The role of artificial intelligence (AI) is growing in almost every area of society. In thinking about the different classes of artificial intelligence technologies that have been developed, one major distinction often made is between “production” technologies and “coordination” technologies (Scott et al., 2007). Production technologies are those that directly help carry out work, either by taking over tasks for humans altogether or by making humans more accurate or efficient in completing tasks. Most technological development begins here, whether we think of early generations of technology (hand tools, machines) or those of the digital age. By contrast, coordination technologies are those that enable the connection or coordination of effort or information from different sources. Examples include algorithms that match people seeking a good or service with the best source to provide it (such as ridesharing or online market platforms). Organizations use coordination technologies to accomplish a variety of work scheduling and task distribution functions. We see additional roles that AI can play in enhancing collective intelligence, from assisting with individual task work to coaching collective teamwork processes (Fiore & Wiltshire, 2016), to diagnosing system-level collaboration breakdowns. We see an opportunity for further development of technologies that can sense and intervene to structure processes proactively before problems emerge. Guided by the transactive system model, we discuss three roles for AI below.

Assistive AI. People are increasingly using what we refer to in Figure 3.4 as “Assistive AI.” These are used to augment all of the cognitive functions mentioned in individual task performance, such as internet search algorithms to aid memory and find information, reminders to direct attention, or even tools that track information in the

Figure 3.4

Potential roles for AI-based team facilitator or coaching agents for improving collective intelligence. Adapted from (Gupta and Woolley, 2021)



environment (e.g. financial indicators, job postings) that inform individual goals. We could also foresee assistive AI that might incorporate knowledge of individual skills or characteristics to specifically augment certain areas of weakness or anticipate reactions that could inhibit judgment and decision making. Such an Assistive AI could arrange the individual's task environment to enhance cognitive functioning. As Stanford University radiologist Curtis Langlotz put it: "AI won't replace radiologists, but radiologists who use AI will replace radiologists who don't." (Dafoe et al., 2021)

Coach AI. Another role for technology is possible in what we refer to as "Coach AI" in enhancing transactive memory, attention, and reasoning systems. The best coaches

structure activities or initiate action proactively rather than waiting for difficulties to develop (Hackman & Wageman, 2005). Some researchers have experimented with relatively simple “nudges” to proactively shape norms or prime awareness of collective memory, attention, or reasoning as a means of sparking the associated transactive processes (Gupta et al., 2019). Newer tools focused on memory sense the content of one member’s work and proactively alert them to related information team members might have (e.g., Gimpel et al., 2020). An important calibration in such systems is to supply the right amount of information, as supplying too many connections could easily lead to overload and defeat the intent of enhanced CI. Additional tools focused on transactive attention include “flash teams” (Retelny et al., 2014) which break complex, multi-faceted projects into segments and coordinate task assignments and hand-offs among contributors in a way that makes it possible to accomplish such projects in significantly less time than standard. In the area of transactive reasoning, some AI-based coach tools facilitate the identification of individual member perspectives on what is most valuable via market-like bidding systems (Malone et al., 2017) to identify the plans or approaches the most members collectively support.

Diagnostic AI. Beyond shaping behavior, there are also possibilities for developing or leveraging AI to monitor a team or larger group via observable collaborative process indicators to diagnose early signs of trouble. Such diagnostics, informed by relationships theorized in this paper, could in turn trigger interventions via Coach AI. For example, monitoring the level of effort and motivation by tracking member activity levels (Engel & Malone, 2018) or emotional states (Van Kleef et al., 2012) can provide early indications of diminishing levels and trigger a review of goal clarity and alignment. Indicators of coordination via activity patterns (Mayo & Woolley, 2021), as well as the concentration of work in fewer projects or areas versus spreading members across many different projects or areas (Bertolotti et al., 2015), can be diagnostic of appropriate performance strategy and

trigger a review of systems for managing attention. Finally, the match between member expertise and time spent on tasks requiring that expertise, along with an increase in specialization among group members, can help diagnose the appropriate use of knowledge and skill (Ren & Argote, 2011). The biggest practical challenge in setting up such systems is in the context-specific identification of operational measures that comprehensively and accurately capture the collaborative indicators.

Human-AI Relationship Management. With the continued development of AI, it is certainly possible to imagine systems that could completely “self-manage” or “self-organize” the interactions of group members to routinely create high levels of collective intelligence. However, we also know that poor human-AI interfacing can complicate what would otherwise be a perfect system on paper (Demir & Cooke, 2014). For instance, work on algorithm aversion demonstrates the numerous situations in which workers override tools and technologies even when doing so to their detriment (Glikson & Woolley, 2020). Systems that would deprive workers of their autonomy to exercise judgment or experience responsibility for their results would undermine their sense of internal motivation (Hackman & Oldham, 1976). Thus it would seem that such systems will need to strike the right balance between autonomy and control, as well as not undermine the natural social cognition that members would normally develop through collaboration (Gupta & Woolley, 2018).

Overall, the systematic development of a Machine Theory of Collective Intelligence holds the promise of producing a virtual coaching system that not only leverages the combination of all three roles discussed in this section but also is capable of building a trustworthy relationship with each human member of the team to successfully influence individual and collective behavior. Such systems would be especially useful in cases where members with varying levels of skills who don't necessarily have experience working with

each other are put together on short notice and are expected to perform under extenuating circumstances (e.g. search and rescue missions in tsunami-affected or war-torn areas).

Conclusion

In developing the relationship between the collaborative processes and transactive systems, this work lays the foundation for translating the Transactive Systems Model of collective intelligence into machine-readable, dynamic indicators that can be used for real-time diagnosis and intervention. Articulation and refinement of context-appropriate metrics hold the potential to provide useful guides for intervention by human managers or AI systems. Together with transactive systems, it pushes the theoretical and practical boundaries of understanding and enhancing human-machine teaming thereby laying a foundation for a Machine Theory of Collective Intelligence.

IV

Summary Conclusions and Future Directions

Modern work has shifted towards increasingly dynamic, temporary, and decentralized forms of organizing as evidenced by the sharp increase in multi-teaming, flash organizations, and open-source work (O’Leary et al., 2011; Valentine et al., 2017). Success in such work leverages new technology for breaking down silos by loosening or disrupting organizational boundaries (Bernstein et al., 2016; Leonardi & Vaast, 2017; Turco, 2016). This represents a significant shift from the central assumptions of the classical administrative theory that values boundaries as a necessary form of structural control needed for efficiently retaining and organizing bounded rational employees (March & Simon, 1993, Simon, 1997). While giving up structural control in the face of dynamism is beneficial for innovation and collaboration (Dougherty & Dunne, 2012; Kellogg et al., 2006), it also results in errors and inefficiencies due to an increased likelihood of worker overload and burnout.

In response, there have been increasing calls for management theory to adapt its assumptions about organizations to more appropriately reflect their reality as complex adaptive systems (Arrow et al., 2000; Cronin et al., 2011; Hackman, 2012). This has resulted in the growing use of concepts originating in research on intelligence in humans and machines (Csaszar & Steinberger, 2021), which has for decades studied the functions that enable systems to adapt and accomplish goals in a wide range of environments that vary in complexity (Legg & Hutter, 2007). And as we see the rapid development of artificial intelligence and its integration into all areas of work, it seems increasingly appropriate to view organizations less as stable structures and more as complex adaptive systems, designed to integrate human intelligence and machine intelligence into collective intelligence.

In this dissertation, I argue that collectives (teams, organizations, or communities) facing highly complex and dynamic situations need to be designed for collective intelligence

and introduce a *Transactive Systems Model of Collective Intelligence*, which provides a process-based framework for understanding how coordination can be accomplished in a complex adaptive system (see Figure 3.1 for an overview of the full model). In a transactive system, members share and allocate cognitive resources and coordinate collective action to their mutual benefit. Instead of using *structural control* to organize and retain employees by administratively bounding organizational situations, this theory proposes that autonomy-seeking members can exercise *transactive control* to engage each other's bounded cognitive capacities in a highly dynamic environment (see Figure 1.1). The work unfolds via agent-based models to support and develop the associated narrative theory and paired with an analysis of archival data to examine similarities and differences between computational models and patterns in the real world.

The key arguments of my dissertation and their contributions are summarized in four points that follow -

1. Collective intelligence is a product of its members' ability to transactively coordinate their *collective attention, memory, and reasoning systems (TAS, TMS, and TRS)*. I describe how each of these transactive systems emerges, bottom-up, from individual-level cognitive processes and member interactions, and identify persistent collective behavioral patterns that indicate their development. In doing so, this work contributes to the literature on social cognition by introducing and extending two new constructs examining collective attention and reasoning (see Chapter 2, and Figure 2.2 for an overview).
2. Collectives exhibit intelligence by continually co-regulating the three transactive systems in order to adapt to their changing environment. The *three-way co-regulation* balances its maintenance function (which "selects the right people and goals") with its efficiency function (which "gets them done quickly"). I theorize that each of these

transactive systems buffers against the negative effects of a specific source of complexity in the environment and their overall influence is thus contingent on the dominant problems experienced by the collective. TMS, TAS, and TRS buffer against knowledge interdependence, workload, and task and resource uncertainty respectively. The integrated theorizing, validated by an agent-based simulation experiment and supported by empirical data from 476 open-source software teams, contributes to the nascent line of research on collective intelligence (see Chapter 2, and Figure 2.4 for an overview).

3. In practice, complex adaptive systems, by their very nature, are extremely difficult to observe and accurately diagnose for the purposes of intervening to improve and maintain collective intelligence. For this, I extend the transactive system model and theorize three observable indicators of collaboration— *level of collective effort*, *appropriateness of performance strategy*, and *use of knowledge and skill*— to be reflective of the ongoing regulatory state of the transactive systems. By theorizing their contribution to efficiency and maintenance functions, I propose that these observable processes can serve as diagnostic indicators and provide real-time information about the non-linear and path-dependent behaviors of the underlying, largely unobservable complex adaptive system. This work extends the research on team coaching providing a theoretical connection between the indicators and underlying socio-cognitive mechanisms and exploring their diagnostic capabilities using agent-based simulation (see Chapter 3, and Figure 3.2 for an overview).
4. Finally, guided by the transactive system model and collaborative indicators, I discuss the potential for developing an “AI-based coaching system,” by describing three distinct roles ripe for AI-driven facilitation— Assistive, Coach, and Diagnostic AI, that can distinctly understand team functioning and successfully nudge individual and

collective behavior in real-time. This pushes the theoretical and practical approach to designing human-machine teams and newer forms of decentralized and autonomous organizations, and calls for the systematic development of a Machine Theory of Collective Intelligence (see Chapter 3, and Figure 3.4 for an overview).

While I've only been able to largely reason about and develop a theoretical model of collective intelligence in this dissertation, it does help us take a helpful perspective towards exploring how existing and newer organizational forms can be categorized and what forms of automation and tooling are likely to be productive in evolving our approach to organizational design. In this final section, I briefly discuss some of these ideas.

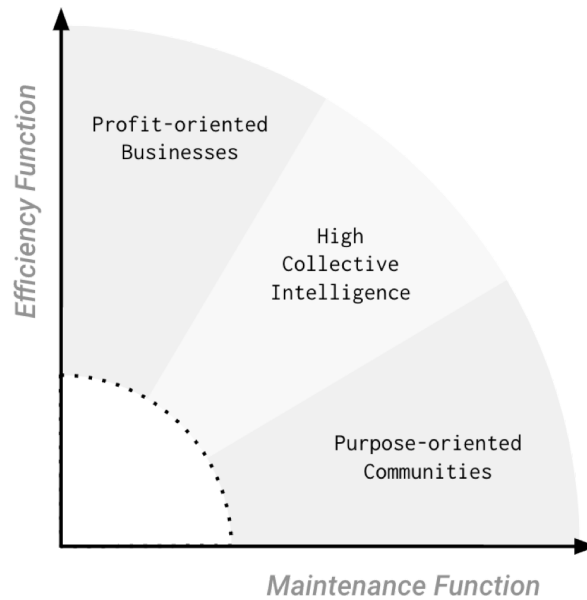
Future Exploration: Developing Organizational Design Principles and Digital Tools for Minimum Viable Collective Intelligence

As discussed in chapter 2, collective intelligence is achieved by successfully fulfilling and balancing its maintenance function (which captures resources by “selecting right people and goals”) with its efficiency function (which deploys said resources by “getting them done quickly”). I anticipate that locating an organization on the *maintenance x efficiency graph* can help its leader and members better understand the prototypical problem they will have to resolve in order to move towards higher collective intelligence (Figure 4.1).

On one end, a typical *profit-oriented business* aims to maximize “technical” efficiency in producing its goods or services. For example, in a software firm, this is achieved by using tools like JIRA that implement the philosophy of Continuous Integration or Continuous Deployment. Such tools standardize the rules for production and coordination of the tasks by providing direct feedback to the developers in short cycles. In a manufacturing firm, this takes the form of varying levels of supply chain automation (like UIPath) to ensure just-in-time materials and inventory management. Thus, profit-oriented businesses can scale

Figure 4.1

Organizational forms located in the Efficiency x Maintenance Graph.



successfully by getting exceedingly good at leveraging technology to reduce and manage their technical debt. Technical debt, the source of operational inefficiencies, is the notion that taking shortcuts while writing software (or making supply chain infrastructure decisions) has undesirable consequences later. This concept of technical debt with an eye for efficiently balancing speed with specialization is explained quite well by Cunningham (1992) in his reflections on building a large portfolio management software -

“Although immature code may work fine and be completely acceptable to the customer, excess quantities will make a program unmasterable, leading to extreme specialization of programmers and finally an inflexible product. Shipping first time code is like going into debt. A little debt speeds development so long as it is paid back promptly with a rewrite. The danger occurs when the debt is not repaid. Every minute spent on not-quite-right code counts as interest on that debt. Entire engineering organizations can be brought to a stand-still under the debt load of an unconsolidated implementation, ...”

On the other end, a typical *purpose-oriented community* aims to maximize its ability to effectively “organize” members from the wider society by creating opportunities for

engaging in activities that tackle problems they personally care about. For example, Wikipedia has created a set of rules that makes it easy for almost anyone to onboard and start contributing to topics that user has knowledge about. These rules and roles continually evolve and are designed with flexibility in mind. They ensure each contribution productively adds to and curates our collective knowledge while allowing for diverse types of members (e.g. varying skills) as well as diverse forms of interactions (e.g. resolving conflicts). In short, a community is successful at pursuing its purpose by dealing with its organizational debt. Organizational debt is generated by “all the people/culture compromises made to ‘just get it done’ in the early stages” (Blank, 2015). Its negative consequences are felt in terms of “the interest companies pay when their structure and policies stay fixed and/or accumulate as the world changes” (Dignan, 2016). Members leave when they feel that the community’s rules (organization structure) are no more serving the purpose that motivated them to join or is inadequate at fulfilling their personal needs. With the uptake of blockchain technology in recent years, we are seeing a surge in experimentation with Decentralized Autonomous Organizations (DAOs) that promise to scale our ability to build and manage communities across the globe. “Instead of a hierarchical structure ... controlling property via the legal system, a DAO involves a set of humans interacting with each other according to a protocol specified in code and enforced on the blockchain” (Nielson, 2021). They are more transparent and have lower barriers to entry than traditional companies (Xie, 2021). DAOs are leveraging software to automate the governance process of negotiating contracts that satisfies the contributors' needs by executing work that furthers the community’s goals.

Seen from the perspective of our collective intelligence model, the profit-oriented business excels in dealing with technical debt by incurring organizational debt that constrains and directs the employees' work, while the purpose-oriented community excels in dealing with organizational debt by incurring technical debt that allows the parallel pursuit of diverse

ideas that don't necessarily build on each other. This sentiment is captured in a recent tweet thread by Sam Spurlin, an org design consultant studying DAOs by becoming a first-hand contributor -

If you want your DAO to run as efficiently as a company, maybe think about starting a company instead. If you want to tap into the nearly limitless potential of humans and create a space for adults to relate as adults in pursuit of a purpose, then go ahead and start a DAO.

DAOs trade efficiency for a huge basket of other great things: resilience, engagement, effectiveness, and innovation to name but a few. If you see inefficiency around you remember what you're "buying" with it.

Not saying efficiency doesn't matter, but if it's the primary thing you're worried about in your DAO it's like going for a swim and spending a lot of effort trying to not get too wet. Kinda comes with the territory and focusing on other things will let you have a much better time. (Spurlin, 2022)

Given the dynamism of our task environment and labor market, we would like to build *high collective intelligence organizations* that are equally adept at navigating and dealing with both kinds of debts. That is, they have rules that are flexible enough to engage everyone who comes up with productive ideas and agile enough to make sure that each contribution adds to or at least complements the existing organizational capabilities. In this setup, the organization manages to have very high efficiency to ensure profitability, and all active employees are motivated to produce their best quality work. A great starting point is to ask: How can a business leverage the wider community without disrupting current operations? How can a community adopt business practices without getting derailed from its mission? And, finally, what specific tools can help us manage this process at scale?

An exciting approach to exploiting organizational forms that hold the promise of exhibiting minimum viable collective intelligence is in the use of experimental platforms like Foundry that enable the deployment of flash teams and flash organizations, "a framework for dynamically assembling and managing paid experts from the crowd" (Retenly et al., 2014, Valentine et al., 2017). This temporary teaming model provides digital tools that scaffold the ability to manage the workflow of short-term decomposable projects. Although, evidence

suggests that this kind of workflow creates intermediate artifacts that inhibit open-ended and creative adaptations typically required for complex design and engineering tasks (Retenly et al., 2017). In other words, flash teams are able to exhibit intelligence when a limited number of people are involved in working on a reasonably well-specified project that can be completed in a short time. But it is unlikely to sustainably manage the debts that accrue when members solve a poorly-specified problem over time. Moreover, scaling the size of a decentralized team comes with its own trade-offs (Hu et al., 2022).

What kind of tools can we imagine that help us, a collective of boundedly rational humans, overcome their fundamental limits to continually balance their technical and organizational debt? Understanding the tradeoffs imposed by algorithmically mediated systems on gig workers is an active area of research that draws scientists and practitioners from computer science and organizational behavior (e.g., Cameron et al., 2021). The DAO tooling community has a fertile ecosystem of open-source software projects that experiment with novel governance and management mechanisms (see Nielson, 2021 for an overview of the common category of infrastructure tools).

As discussed at the end of the previous chapter, the Transactive Systems Model along with the indicators has the potential for informing the development of an “*Artificial Social Intelligence*” that can passively read digital trace data generated by member interactions in real-time and deliver dashboards and APIs (data and program interfaces) that help diagnose the current state of the collective as well as generate recommendations for individual and joint intervention. Moreover, such a system can operate at multiple scales simultaneously (teams, nested in communities, nested across organizations). I am involved in prototyping such a real-time coaching system for teams playing dynamically perturbed search and rescue missions in Minecraft (ASIST Grant, 2019-22).

Although such tools help in the day-to-day management of collectives, a fully fleshed-out Machine Theory of Collective Intelligence will also be able to produce tools that help us explore the high dimensional design space for collectives. Just as modern software packages help engineers design bridges or test the performance of aircraft designs by modeling air flows, a software for *Computer-Aided Socio-Technical Systems Design* will make it easy to prototype collectives and simulate the effects of different rules for the management of technical and organizational debts. A well-designed user interface will make it easy for anyone to thoughtfully select the features of humans and machines that make up the collective as well as systematically model the rules governing the interactions for each combination.

Managerial Implications

While the goal of this dissertation was to develop a holistic theory about how complex adaptive systems exhibit collective intelligence, the transactive systems model lens generates two major insights for managers dealing with challenges posed by dynamic coordination. First, if managers fail to note the need to respond adaptively to the dynamism in their (and their team members) work environment they will likely never solve the correct problem. As structural control mechanisms have worked quite well in the traditional work settings, the temptation to double-down on them as the work gets more dynamic due to changes in coordination technology is really high. The premise of our theory suggests that doing so may cause more harm. The model recommends that in dynamic environments, it is useful for the managers to lean into it and get better at dynamic coordination by further loosening structures and giving their members more autonomy. This does come at a cost. Making various organizational boundaries more permeable will increase both the informational and relational load experienced by the members. In contemporary teams and organizations, the functional

skills members have are considered their biggest asset. With the loosening of boundaries, it becomes easy for them to stumble upon opportunities outside their present organization that find their skills highly desirable. Furthermore, the rise of remote and flexible work setups has drastically reduced the cost of changing employers and moving elsewhere. Without strong ties within the team that reinforces their value in the current role, highly skilled employees are bound to be lost in the constant presence of alternatives. The emphasis of the manager's job is thus shifting from efficient project management to employee retention. In such circumstances, the role of the manager shifts from exercising structural control by reinforcing boundaries to facilitating transactive control by aligning and building shared mental models around collective goals, needs, availabilities, and skills. Specifically, as per the Transactive Systems Model, work is accomplished when collaborators are able to reciprocally rely on each other by better understanding their skills, availabilities, and goals. Hence, the manager's function is to ensure that this mutual understanding doesn't break down.

Boiling it down, three desirable characteristics of a manager of dynamic collectives are: (1) They consistently communicate and enforce a few expectations - Every team needs ongoing goals, norms, and values. The manager is skilled in identifying a few that are critical for the collective's goals and enforcing them by articulating clearly, reminding regularly, as well as publicly recognizing actions that align, and privately correcting those that misalign. Regular reinforcement makes it easier to both maintain the collective's transactive reasoning system and detect when the expectations need adjustment. (2) The manager employs their own expertise as a multiplier, and not a limiter - In a changing environment, the management is not in the best position to understand all the relevant influences underlying how the situation has changed. Rather than enforcing their point of view, their expertise can be better utilized to elicit the employees' thinking and facilitate collective integration and decision making. Facilitation ensures the emergence and adoption of new shared norms that suit the

new situation faced by the collective thereby rebalancing across the three transactive systems instead of getting stuck in any one at the neglect of others. (3) The manager doesn't direct but unlocks and unleashes the employees' potential - They are skilled in facilitating transactive interactions across many people they manage. Building a better understanding of their team members' skills, workloads, and ambitions makes it easier for the manager to tap into organizational resources that are likely to generate win-win scenarios for both, their team and the larger organization. Ultimately, rather than molding each employee into a set role, or coaching them along a pre-defined career trajectory, such a manager spends time getting to know the evolving goals and motives of the people they manage and creating opportunities that help them become the best versions of who they are.

Finally, with the emphasis on transactive forms of coordination, the ideal member of a high intelligence collective is what one would deem to be a "super-collaborator." In addition to being highly skilled in their area of expertise, they are also highly metacognitive. For this reason, managers should focus on evaluating the collaboration skills of the people they manage in addition to their technical skills. Specifically, a self and peer assessment of how well-developed are the focal member's meta-memory, meta-attention, and meta-reasoning processes. For example, when given a list of common tasks, all members of a team with a well-developed meta-attention will independently report a similar rank-ordering for triaging through the task list. Similarly, when given a commonly faced problem, all members of a team with a well-developed meta-memory will independently report a similar subset of members they will reach out for help and resources that they will access to resolve the problem at hand. Members who exhibit a high level of metacognition across all three functions will be exceptionally good at dynamically understanding and engaging in work with their collaborators. Lastly, the managers themselves should also be evaluated on their ability to monitor, diagnose and coach the collaboration of their team. In addition to the

previously discussed characteristics, good managers will have a firm grasp of the extent to which the people they manage are motivated to put in their effort, their collective effort is appropriately deployed and the members are not only using their skills but also have an opportunity to grow and develop more skills that can be brought to bear on collective tasks in the future.

Conclusion

As articulated in the introduction the work presented here is inspired by and seeks to build on the work of Herbert Simon, who proposed that the boundedly-rational *administrative man* (sic) organizes others via structural control to efficiently absorb uncertainty and fulfill collective goals (March & Simon, 1993, Simon, 1997). My work in this dissertation builds upon Simon's main thesis and develops a path that I hope will expand the scope and applicability of his ideas in the future. I propose that when dynamism is very high, the boundedly-rational *metacognitive individual*, being aware of their own limitations, cooperates with others via transactive interactions to manage uncertainty and satisfactorily balance the fulfillment of personal and collective goals. Metacognition enables transactive behavior and the emergent collective attention, memory, and reasoning structures necessary for the development of collective intelligence. While much remains to be learned about the nature of collective intelligence in complex adaptive systems, this dissertation introduces a socio-cognitive model to enable a multi-method systems approach for investigating the emergence, adaptation, and diagnosis of collective intelligence, laying the groundwork for future research. In so doing, it is my hope that this work embodies the qualities articulated by the researcher who is my strongest source of intellectual inspiration—

“The goal of science is to make the wonderful and the complex understandable and simple - but not less wonderful.”

- Herbert Simon, *The Sciences of the Artificial* (1998)

Appendix

Table A1

Procedure for Computational Model for Emergence of Collective Intelligence.

Step 1: Clean up Short- and Long-term Memory (assume negligible time)

- 1.1: Memory Maintenance: Forget Old Experiences
- 1.2: Attention Maintenance: Unload knowledge that is not recently used

Step 2: Evaluate and Store all MetaKnowledge Messages (assume negligible time)

- 2.1: Meta-memory: Directory Updating
- 2.2: Meta-attention: Availability Updating
- 2.3: Meta-reasoning: Goals and Needs Updating

Step 3: Evaluate Inbox and Select Focal Task (assume negligible time)

- 3.1: Triage my inbox: Estimate the best candidate for each task as per appropriate rule-set (Table A2)
 - 3.1.1: Deploy appropriate rule-set using my meta-attention (workload), meta-memory (skill), and meta-reasoning (needs) to assign tasks to the best candidate.
 - 3.1.2: If my current workload (Level-Change); send meta-attention message to all members
- 3.2: Select focal task for work:
 - 3.2.1: If inbox is empty: Do nothing; END timestep
 - 3.2.2: If TMS (ON): Select the oldest task
 - 3.2.3: If TAS (ON): Select the highest priority task
 - 3.2.4: If Needs (ON and unfulfilled): Select task that best suits personal goals

Step 4: Load Focal Task in Attention (top-down selection of knowledge)

- 4.1: Figure out the current task-state and identify the skill knowledge it needs
- 4.2: Load skill knowledge in attention from memory
 - 4.2.1: If knowledge already in attention, continue to step 4
 - 4.2.2: Else spend timestep loading knowledge (switching cost) and END timestep

Step 5: Work on Focal Task, Learn from it and Regulate motives and personal goals

- 5.1: Calculate Skill multiplier for work generation based on level of relative task difficulty (default = 1)
 - 5.1.1: If too easy (< my level - 200); multiplier = 2.0
 - 5.1.2: If too hard (> my level + 200); multiplier = 0.5
- 5.2: Calculate Effort multiplier for work generation based on current motivational state (default = 1)
 - 5.2.1: If MemberNeeds (ON) and Motivational Need (Below Threshold); multiplier = 0.5
- 5.3: Generate Work proportional to my skill-level, learn, and regulate personal goals
 - 5.3.1: $WorkGenerated = Skill\text{-}level * Skill\text{-}multiplier * Effort\text{-}multiplier$
 - 5.3.2: Learning by Experience: increase experience-level proportional to learning rate
 - 5.3.2.1: If Skill-knowledge (Level-Up); send meta-memory message to all members
 - 5.3.2: Goal Regulation: update needs satisfaction levels based on current task and skill
 - 5.3.2.1: If Needs (Level-Change); send meta-reasoning message to all members
- 5.4: Execute task
 - 5.4.1: Spend timestep working for the identified skill ($Taskload = Taskload - WorkGenerated$)
 - 5.4.2: If task (Complete): Remove task from inbox and END timestep

Note. All steps of the pseudo-code were evaluated individually by each agent for each timestep. But the order in which each member evaluated these steps was randomized for each timestep.

Table A2

Rulesets for Triaging a Member's Inbox.

Ruleset 1: If TMS (OFF) and TAS (OFF): Do nothing

Ruleset 2: If TMS (ON) and TAS (OFF): Ask meta-memory: Who is the 'skill' expert?

(NOT ME): Send the task to the expert

(ME or NO ANSWER): Keep the task

Ruleset 3: If TMS (OFF) and TAS (ON): Ask meta-attention: Is this my focal task?

(YES): Keep the task

(NO): Ask meta-attention: Who is working on a lower 'priority' focal task?

(NOT ME): Send the task to the available member

(ME or NO ANSWER): Keep the task

Ruleset 4: If TMS (ON) and TAS (ON): Ask meta-memory: Am I the 'skill' expert?

(YES): Ask meta-attention: Is this my focal task?

(YES): Keep the task

(NO): Ask meta-memory & loop through next best 'skill' expert:

Ask meta-attention: Is this member working on a lower 'priority' task?

(YES): Ask for help from available non-expert; 'Send Task'

(NO ANSWER): Keep the task

(NO): Ask meta-memory & loop through all relative 'skill' experts:

Ask meta-attention: Is this member working on a lower 'priority' task?

(YES): Ask for help from the available expert; 'Send Task'

(NO ANSWER): Keep the task

Ruleset 5: If TRS (ON): Before 'Send Task' in Ruleset 4: Ask meta-reasoning: Are receivers' needs satisfied?

(YES) Send task

(NO) Ask meta-reasoning: Can this task align with receivers' unsatisfied needs?

(YES) Provide goal support and send task

(NO) Keep the task

Note. Members iterate through each task in their inbox and figure out task-state (identify skill needed and priority). Then estimates the best candidate for each task based on the appropriate rule-set.

References

- Ackerman, R., & Thompson, V. A. (2017). Meta-Reasoning: Monitoring and Control of Thinking and Reasoning. *Trends in Cognitive Sciences*, 21(8), 607–617. <https://doi.org/10.1016/j.tics.2017.05.004>
- Afuah, A., & Tucci, C. L. (2012). Crowdsourcing As a Solution to Distant Search. *Academy of Management Review*, 37(3), 355–375. <https://doi.org/10.5465/amr.2010.0146>
- Ancona, D. G. & Caldwell, D. F. (1992). Bridging the boundary: External activity and performance in organizational teams. *Administrative Science Quarterly*, 37(4): 634–665.
- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Lawrence Erlbaum Associates.
- Argote, L., & Greve, H. R. (2007). “A Behavioral Theory of the Firm”: 40 Years and Counting: Introduction and Impact. *Organization Science*, 18(3), 337–349.
- Arrow, H., McGrath, J. E., & Berdahl, J. L. (2000). *Small groups as complex systems: Formation, coordination, development and adaptation*. Sage Publications.
- Bacharach, M. (1999). Interactive team reasoning: A contribution to the theory of co-operation. *Research in Economics*, 53(2), 117–147. <https://doi.org/10.1006/reec.1999.0188>
- Bakker, R. M. (2010). Taking stock of temporary organizational forms: A systematic review and research agenda. *International Journal of Management Reviews*, 12(4): 466–486.
- Bandura, A. (1991). Social cognitive theory of self-regulation. *Organizational Behavior and Human Decision Processes*, 50(2), 248–287. [https://doi.org/10.1016/0749-5978\(91\)90022-L](https://doi.org/10.1016/0749-5978(91)90022-L)
- Baron-Cohen, S. (2000). Theory of mind and autism: A review. In *International Review of Research in Mental Retardation* (Vol. 23, pp. 169–184). Academic Press. [https://doi.org/10.1016/S0074-7750\(00\)80010-5](https://doi.org/10.1016/S0074-7750(00)80010-5)
- Bechky, B. A. (2003). Sharing meaning across occupational communities: The transformation of understanding on a production floor. *Organization Science*, 14(3): 312–330.
- Bernstein, E. S., Bunch, J., Canner, N. & Lee, M. (2016). Beyond the holacracy hype. *Harvard Business Review*, 94(7): 38–49.
- Bernstein, E., Shore, J., & Lazer, D. (2018). How intermittent breaks in interaction improve collective intelligence. *Proceedings of the National Academy of Sciences*, 115(35), 8734–8739. <https://doi.org/10.1073/pnas.1802407115>
- Bertolotti, F., Mattarelli, E., Vignoli, M., & Macri, D. M. (2015). Exploring the relationship between multiple team membership and team performance: The role of social networks and collaborative technology. *Research Policy*, 44(4), 911–924. <https://doi.org/10.1016/j.respol.2015.01.019>
- Blank, S. (2015). *Organizational Debt is like Technical debt – but worse*. <https://steveblank.com/2015/05/19/organizational-debt-is-like-technical-debt-but-worse/>
- Buckley, W. (1967). *Sociology and modern systems theory*. Prentice-Hall.
- Burke, C. S., Stagl, K. C., Salas, E., Pierce, L., & Kendall, D. (2006). Understanding team adaptation: A conceptual analysis and model. *The Journal of Applied Psychology*, 91(6), 1189–1207. <https://doi.org/10.1037/0021-9010.91.6.1189>
- Cameron, L., Christin, A., Ann DeVito, M., R. Dillahunt, T., Elish, M., Gray, M., Qadri, R., Raval, N., Valentine, M., & Anne Watkins, E. (2021). “This Seems to Work”: Designing Technological Systems with The Algorithmic Imaginations of Those Who Labor. *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–5. <https://doi.org/10.1145/3411763.3441331>

- Cappelli, P. (2020). Stop Overengineering People Management. *Harvard Business Review*.
<https://hbr.org/2020/09/stop-overengineering-people-management>
- Carley, K. M. (2001). Computational Approaches to Sociological Theorizing. In J. H. Turner (Ed.), *Handbook of Sociological Theory* (pp. 69–83). Springer US.
https://doi.org/10.1007/0-387-36274-6_4
- Carsrud, A., Brännback, M., Elfving, J., & Brandt, K. (2009). Motivations: The Entrepreneurial Mind and Behavior. In A. L. Carsrud & M. Brännback (Eds.), *Understanding the Entrepreneurial Mind: Opening the Black Box* (pp. 141–165). Springer.
- Cohen, M. D. (2007). “Administrative Behavior”: Laying the Foundations for Cyert and March. *Organization Science*, 18(3), 503–506.
- Cronin, M. A., Weingart, L. R., & Todorova, G. (2011). Dynamics in Groups: Are We There Yet? *Academy of Management Annals*, 5(1), 571–612.
<https://doi.org/10.5465/19416520.2011.590297>
- Csaszar, F., & Steinberger, T. (2021). Organizations as artificial intelligences: The use of artificial intelligence analogies in organization theory. *Academy of Management Annals*. <https://doi.org/10.5465/annals.2020.0192>
- Cummings, J. N. & Haas, M. R. (2012). So many teams, so little time: Time allocation matters in geographically dispersed teams. *Journal of Organizational Behavior*, 33(3): 316–341.
- Cunningham, W. (1992). *OOPSLA Experience Report: The WyCash Portfolio Management System*. Retrieved April 11, 2022, from <http://c2.com/doc/oopsla92.html>
- Dafoe, A., Bachrach, Y., Hadfield, G., Horvitz, E., Larson, K., & Graepel, T. (2021). Cooperative AI: Machines must learn to find common ground. *Nature*, 593(7857), 33–36. <https://doi.org/10.1038/d41586-021-01170-0>
- 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.
- Demetriou, A. (2000). Organization and development of self-understanding and self-regulation: Toward a general theory. In *Handbook of self-regulation* (pp. 209–251). Academic Press. <https://doi.org/10.1016/B978-012109890-2/50036-6>
- Demir, M., & Cooke, N. J. (2014). Human Teaming Changes Driven by Expectations of a Synthetic Teammate. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 58(1), 16–20.
- Desimone, R., & Duncan, J. (1995). Neural Mechanisms of Selective Visual Attention. *Annual Review of Neuroscience*, 18(1), 193–222.
<https://doi.org/10.1146/annurev.ne.18.030195.001205>
- Deutschkron, S., & Pearce, C. (2019). *Sixth annual “Freelancing in America” study finds that more people than ever see freelancing as a long-term career path | Upwork*.
<https://www.upwork.com/press/releases/freelancing-in-america-2019>
- Dignan, A. (2022, April 6). How To Eliminate Organizational Debt. *The Ready*.
<https://medium.com/the-ready/how-to-eliminate-organizational-debt-8a949c06b61b>
- Dolan, R. J. (2002). Emotion, Cognition, and Behavior. *Science*, 298(5596), 1191–1194.
<https://doi.org/10.1126/science.1076358>
- Dougherty, D. & Dunne, D. D. (2012). Digital science and knowledge boundaries in complex innovation. *Organization Science*, 23(5): 1467–1484.
- Druskat, V. U., & Wolff, S. B. (2001). Building the Emotional Intelligence of Groups. *Harvard Business Review*.
<https://hbr.org/2001/03/building-the-emotional-intelligence-of-groups>
- Edmondson, A. C., & Harvey, J.-F. (2018). Cross-boundary teaming for innovation: Integrating research on teams and knowledge in organizations. *Human Resource*

- Management Review*, 28(4), 347–360. <https://doi.org/10.1016/j.hrmr.2017.03.002>
- Egeth, H. E., & Yantis, S. (1997). Visual attention: Control, representation, and time course. *Annual Review of Psychology*, 48, 269–297. <https://doi.org/10.1146/annurev.psych.48.1.269>
- Eisenman, M., Paruchuri, S., & Puranam, P. (2020). The design of emergence in organizations. *Journal of Organization Design*, 9(1), 25. <https://doi.org/10.1186/s41469-020-00089-6>
- Endsley, M. R. (2015). Situation Awareness Misconceptions and Misunderstandings. *Journal of Cognitive Engineering and Decision Making*, 9(1), 4–32. <https://doi.org/10.1177/1555343415572631>
- Engel, D., & Malone, T. W. (2018). Integrated information as a metric for group interaction. *PLoS One*, 13(10), e0205335.
- Engel, D., Woolley, A. W., Jing, L. X., Chabris, C. F., & Malone, T. W. (2014). Reading the Mind in the Eyes or Reading between the Lines? Theory of Mind Predicts Collective Intelligence Equally Well Online and Face-To-Face. *PLoS ONE*, 9(12), e115212. <https://doi.org/10.1371/journal.pone.0115212>
- Epstein, J. M. (1999). Agent-based computational models and generative social science. *Complexity*, 4(5), 41–60.
- Faraj, S. & Yan, A. (2009). Boundary work in knowledge teams. *Journal of Applied Psychology*, 94(3): 604–617.
- Fiore, S. M., & Wiltshire, T. J. (2016). Technology as Teammate: Examining the Role of External Cognition in Support of Team Cognitive Processes. *Frontiers in Psychology*, 7. <https://doi.org/10.3389/fpsyg.2016.01531>
- Fitzsimons, G. M., Sackett, E., & Finkel, E. J. (2016). Transactive Goal Dynamics Theory: A relational goals perspective on work teams and leadership. *Research in Organizational Behavior*, 36, 135–155. <https://doi.org/10.1016/j.riob.2016.11.006>
- Flavell, J. H., & Wellman, H. M. (1977). Metamemory. In R. V. Kail & J. W. Hagen (Eds.), *Perspectives on the development of memory and cognition*. Lawrence Erlbaum.
- Gimpel, H., Graf-Drasch, V., Laubacher, R. J., & Wöhl, M. (2020). Facilitating like Darwin: Supporting cross-fertilisation in crowdsourcing. *Decision Support Systems*, 132, 113282.
- Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2), 627–660. <https://doi.org/10.5465/annals.2018.0057>
- Grand, J. A., Braun, M. T., Kuljanin, G., Kozlowski, S. W. J., & Chao, G. T. (2016). The dynamics of team cognition: A process-oriented theory of knowledge emergence in teams. *The Journal of Applied Psychology*, 101(10), 1353–1385. <https://doi.org/10.1037/apl0000136>
- Gupta, P., Kim, Y. J., Glikson, E., and Woolley, A. W. (Under Review). Digitally Nudging Collaborative Processes to Enhance Collective Intelligence.
- Gupta, P., Kim, Y. J., Glikson, E., & Woolley, A. W. (2019). Digitally nudging team processes to enhance collective intelligence. *Proceedings of Collective Intelligence 2019*.
- Gupta, P., & Woolley, A. W. (2018). Productivity in an era of multi-teaming: The role of information dashboards and shared cognition in team performance. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW), 62:1-62:18. <https://doi.org/10.1145/3274331>
- Hackman, J. R. (1987). The design of work teams. In J. Lorsch (Ed.), *Handbook of organizational behavior*. Prentice-Hall.
- Hackman, J. R. (2011). *Collaborative Intelligence: Using Teams to Solve Hard Problems*. Berrett-Koehler Publishers.

- Hackman, J. R. (2012). From causes to conditions in group research. *Journal of Organizational Behavior*, 33(3), 428–444. <https://doi.org/10.1002/job.1774>
- Hackman, J. R., & Oldham, G. R. (1976). Motivation through the design of work: Test of a theory. *Organizational Behavior and Human Decision Processes*, 16(2), 250–279.
- Hackman, J. R. & Wageman, R. (1995). Total quality management: Empirical, conceptual, and practical issues. *Administrative Science Quarterly*, 40(2): 309–342.
- Hackman, J. R., & Wageman, R. (2005). A theory of team coaching. *Academy of Management Review*, 30(2), 269–287. <https://doi.org/10.5465/amr.2005.16387885>
- Hackman, J. R., & Wageman, R. (2009). Foster Team Effectiveness by Fulfilling Key Leadership Functions. In *Handbook of Principles of Organizational Behavior* (pp. 273–293). John Wiley & Sons, Ltd. <https://doi.org/10.1002/9781119206422.ch15>
- Hong, L., & Page, S. E. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences*, 101(46), 16385–16389. <https://doi.org/10.1073/pnas.0403723101>
- Hu, X. E., Hinds, R., Valentine, M. A., & Bernstein, M. S. (2022). A “Distance Matters” Paradox: Facilitating Intra-Team Collaboration Can Harm Inter-Team Collaboration. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW1), 1–36. <https://doi.org/10.1145/351289>
- Kahneman, D. (1973). *Attention and effort*. Prentice-Hall.
- Kellogg, K. C., Orlikowski, W. J. & Yates, J. (2006). Life in the trading zone: Structuring coordination across boundaries in postbureaucratic organizations. *Organization Science*, 17(1): 22–44.
- Kim, Y. J., Engel, D., Woolley, A. W., Lin, J., McArthur, N., & Malone, T. W. (2017). What makes a strong team? Using collective intelligence to predict performance of teams in League of Legends. *Proceedings of ACM CSCW 2017*.
- Knudsen, E. I. (2007). Fundamental Components of Attention. *Annual Review of Neuroscience*, 30(1), 57–78. <https://doi.org/10.1146/annurev.neuro.30.051606.094256>
- Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: Towards the underlying neural circuitry. *Human Neurobiology*, 4(4), 219–227.
- Koestner, R., Powers, T. A., Carbonneau, N., Milyavskaya, M., & Chua, S. N. (2012). Distinguishing Autonomous and Directive Forms of Goal Support: Their Effects on Goal Progress, Relationship Quality, and Subjective Well-Being. *Personality and Social Psychology Bulletin*, 38(12), 1609–1620. <https://doi.org/10.1177/0146167212457075>
- Kozlowski, S. W. J., Chao, G. T., Grand, J. A., Braun, M. T., & Kuljanin, G. (2013). Advancing Multilevel Research Design: Capturing the Dynamics of Emergence. *Organizational Research Methods*, 16(4), 581–615. <https://doi.org/10.1177/1094428113493119>
- Kozlowski, S. W. J., & Klein, K. J. (2000). A multilevel approach to theory and research in organizations: Contextual, temporal, and emergent processes. In *Multilevel theory, research, and methods in organizations: Foundations, extensions, and new directions* (pp. 3–90). Jossey-Bass.
- Krackhardt D. M (2000). Modeling structures of organizations. In Ilgen, D.R. & Hulin, C.L. (Eds.), *Computational modeling of behavior in organizations: The third scientific discipline*, Quorum Books, Westport, CT, pp. 269-274
- Laird, J. (2012). *The Soar cognitive architecture*. MIT Press.
- Lavie, N., Hirst, A., de Fockert, J. W., & Viding, E. (2004). Load theory of selective attention and cognitive control. *Journal of Experimental Psychology. General*, 133(3), 339–354. <https://doi.org/10.1037/0096-3445.133.3.339>
- Lee, M. Y. & Edmondson, A. C. (2017). Self-managing organizations: Exploring the limits of

- less-hierarchical organizing. *Research in organizational behavior*, 37: 35–58.
- Legg, S., & Hutter, M. (2007). Universal Intelligence: A Definition of Machine Intelligence. *Minds and Machines*, 17(4), 391–444. <https://doi.org/10.1007/s11023-007-9079-x>
- Leonardi, P. M. & Vaast, E. (2017). Social media and their affordances for organizing: A review and agenda for research. *Academy of Management Annals*, 11(1): 150–188.
- Leonard, Naomi E., and Levin, Simon A. (2022; in press) Collective intelligence as a public good. *Collective Intelligence*.
- Lewis, K. (2004). Knowledge and Performance in Knowledge-Worker Teams: A Longitudinal Study of Transactive Memory Systems. *Management Science*, 50(11), 1519–1533. <https://doi.org/10.1287/mnsc.1040.0257>
- Lichtman, G. (2018). Increasing an NPM package’s search score. Retrieved from: <https://itnext.io/increasing-an-npm-packages-search-score-fb557f859300>
- Luria, A. R. (1973). *The working brain: An introduction to neuropsychology*. Basic Books.
- Malone, T. W. (2018). *Superminds: The surprising power of people and computers thinking together* (First edition). Little, Brown and Company.
- Malone, T. W., & Bernstein, M. S. (Eds.). (2015). *Handbook of collective intelligence*. MIT Press.
- Malone, T. W., Nickerson, J. V., Laubacher, R. J., Fisher, L. H., De Boer, P., Han, Y., & Towne, W. B. (2017). Putting the pieces back together again: Contest webs for large-scale problem solving. *Proceedings of ACM CSCW 2017*, 1661–1674.
- March, J. G. (1991). Exploration and Exploitation in Organizational Learning. *Organization Science*, 2(1), 71–87.
- March, J. G. (2006). Rationality, foolishness, and adaptive intelligence. *Strategic Management Journal*, 27(3), 201–214. <https://doi.org/10.1002/smj.515>
- March, J. G., & Olsen, J. P. (1979). *Ambiguity and choice in organizations* (2d ed). Universitetsforlaget.
- March, J. G., & Shapira, Z. (1992). Variable risk preferences and the focus of attention. *Psychological Review*, 99(1), 172–183. <https://doi.org/10.1037/0033-295X.99.1.172>
- March, J. G., Simon, H. A., & Guetzkow, H. S. (1993). *Organizations* (2nd ed). Blackwell.
- Mathieu, J. E., Gallagher, P. T., Domingo, M. A., & Klock, E. A. (2019). Embracing Complexity: Reviewing the Past Decade of Team Effectiveness Research. *Annual Review of Organizational Psychology and Organizational Behavior*, 6(1), 17–46.
- Mayo, A.T. (in press). Syncing up: A process model of emergent interdependence in dynamic teams. *Administrative Science Quarterly*.
- Mayo, A. T., & Woolley, A. W. (2021). Variance in group ability to transform resources into performance, and the role of coordinated attention. *Academy of Management Discoveries*, 7(2), 225–246. <https://doi.org/10.5465/amd.2019.0231>
- Meadows, D. H., & Wright, D. (2008). *Thinking in systems: A primer*. Chelsea Green Pub.
- Meslec, N., Aggarwal, I., & Curseu, P. L. (2016). The Insensitive Ruins It All: Compositional and Compilational Influences of Social Sensitivity on Collective Intelligence in Groups. *Frontiers in Psychology*, 7. <https://doi.org/10.3389/fpsyg.2016.00676>
- Miller, J. H., & Page, S. E. (2007). *Complex adaptive systems: An introduction to computational models of social life*. Princeton University Press.
- Miller, P. H., & Bigi, L. (1979). The development of children’s understanding of attention. *Merrill-Palmer Quarterly*, 25(4), 235–250.
- Moreland, R. L., Levine, J. M., & Cini, M. A. (1993). Group socialization: The role of commitment. In *Group motivation: Social psychological perspectives* (pp. 105–129). Harvester Wheatsheaf.
- Mortensen, M. (2014). Constructing the team: The antecedents and effects of membership model divergence. *Organization Science*, 25(3): 909–931.

- Mortensen, M., & Haas, M. R. (2018). Perspective—Rethinking Teams: From Bounded Membership to Dynamic Participation. *Organization Science*, 29(2), 341–355. <https://doi.org/10.1287/orsc.2017.1198>
- Neeley, T. B. & Leonardi, P. M. (2018). Enacting knowledge strategy through social media: Passable trust and the paradox of nonwork interactions. *Strategic Management Journal*, 39(3): 922–946.
- Newell, A. (1990). *Unified theories of cognition*. Harvard University Press.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Prentice-Hall.
- Nielsen, K. (2021, April 23). *What is a DAO? Beginners Guide to DAO Ecosystems*. The Defiant. <https://thedefiant.io/what-is-a-dao/>
- Ocasio, W. (2011). Attention to attention. *Organization Science*, 22(5), 1286–1296.
- O’Leary, M. B., & Cummings, J. N. (2007). *The Spatial, Temporal, and Configurational Characteristics of Geographic Dispersion in Teams* (SSRN Scholarly Paper ID 1739905). Social Science Research Network. <https://papers.ssrn.com/abstract=1739905>
- O’Leary, M. B., Mortensen, M., & Woolley, A. W. (2011). Multiple Team Membership: A Theoretical Model of its Effects on Productivity and Learning for Individuals and Teams. *Academy of Management Review*, 36(3), 461–478. <https://doi.org/10.5465/AMR.2011.61031807>
- O’Leary, M. B., Woolley, A. W., & Mortensen, M. (2012). Multiteam membership in relation to multiteam systems. In *Multiteam systems: An organization form for dynamic and complex environments* (pp. 141–172). Routledge/Taylor & Francis Group.
- Page, S. E. (2006). Path Dependence. *Quarterly Journal of Political Science*, 1(1), 87–115. <https://doi.org/10.1561/100.00000006>
- Raveendran, M., Silvestri, L., & Gulati, R. (2020). The Role of Interdependence in the Micro-Foundations of Organization Design: Task, Goal, and Knowledge Interdependence. *Academy of Management Annals*, 14(2), 828–868. <https://doi.org/10.5465/annals.2018.0015>
- Ren, Y., & Argote, L. (2011). Transactive memory systems 1985–2010: An integrative framework of key dimensions, antecedents, and consequences. *The Academy of Management Annals*, 5(1), 189–229. <https://doi.org/10.1080/19416520.2011.590300>
- Retelny, D., Bernstein, M. S., & Valentine, M. A. (2017). No Workflow Can Ever Be Enough: How Crowdsourcing Workflows Constrain Complex Work. *Proceedings of the ACM on Human-Computer Interaction*, 1(CSCW), 1–23. <https://doi.org/10.1145/3134724>
- Retelny, D., Robaszkiewicz, S., To, A., Lasecki, W. S., Patel, J., Rahmati, N., Doshi, T., Valentine, M., & Bernstein, M. S. (2014). Expert crowdsourcing with flash teams. *Proceedings of ACM Symposium on User Interface Software and Technology*, 75–85.
- Rico, R., Gibson, C. B., Sánchez-Manzanares, M., & Clark, M. A. (2019). Building team effectiveness through adaptation: Team knowledge and implicit and explicit coordination. *Organizational Psychology Review*, 9(2–3), 71–98. <https://doi.org/10.1177/2041386619869972>
- Riedl, C., Kim, Y. J., Gupta, P., Malone, T. W., & Woolley, A. W. (2021). Quantifying collective intelligence in human groups. *Proceedings of the National Academy of Sciences*, 118(21). <https://doi.org/10.1073/pnas.2005737118>
- Riedl, C., & Woolley, A. W. (2017). Teams vs. crowds: A field test of the relative contribution of incentives, member ability, and emergent collaboration to crowd-based problem solving performance. *Academy of Management Discoveries*, 3(4), 382–403.
- Ryan, R. M., & Deci, E. L. (2018). *Self-Determination Theory: Basic Psychological Needs in Motivation, Development, and Wellness* (1 edition). The Guilford Press.

- Saner, L. D., Bolstad, C. A., Gonzalez, C., & Cuevas, H. M. (2009). Measuring and Predicting Shared Situation Awareness in Teams. *Journal of Cognitive Engineering and Decision Making*, 3(3), 280–308. <https://doi.org/10.1518/155534309X474497>
- Saunders, C. S. & Ahuja, M. K. (2006). Are all distributed teams the same? Differentiating between temporary and ongoing distributed teams. *Small Group Research*, 37(6): 662–700.
- Schwaninger M. & Groesser S. (2009) System Dynamics Modeling: Validation for Quality Assurance. In: Meyers R. (eds) *Complex Systems in Finance and Econometrics*. Springer, New York, NY. https://doi.org/10.1007/978-1-4419-7701-4_42
- Schwarz, N. (2000). Emotion, cognition, and decision making. *Cognition & Emotion*, 14(4), 433–440.
- Scott, W. R., Davis, G. F., & Scott, W. R. (2007). *Organizations and organizing: Rational, natural, and open system perspectives* (1st ed). Pearson Prentice Hall.
- Simon, H. A. (1973). Applying Information Technology to Organization Design. *Public Administration Review*, 33(3), 268–278. <https://doi.org/10.2307/974804>
- Simon, H. A. (1997). *Administrative Behavior: A study of decision-making processes in administrative organizations* (4th ed). Free Press.
- Spurlin, S. [@samspurlin]. (2022, April 5). *If you want your DAO to run as efficiently* [Tweet]. Twitter. <https://twitter.com/samspurlin/status/1511300823938129924>
- Taylor, F. W. (2012). *The Principles of Scientific Management*. The Floating Press.
- Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *The Behavioral and Brain Sciences*, 28(5), 675–691; discussion 691-735. <https://doi.org/10.1017/S0140525X05000129>
- Turco, C. J. (2016). *The conversational firm: Rethinking bureaucracy in the age of social media*. New York: Columbia University Press.
- Valentine, M. A., & Edmondson, A. C. (2015). Team scaffolds: How mesolevel structures enable role-based coordination in temporary groups. *Organization Science*, 26(2), 405–422. <https://doi.org/10.1287/orsc.2014.0947>
- Valentine, M. A., Retelny, D., To, A., Rahmati, N., Doshi, T., & Bernstein, M. S. (2017). Flash organizations: Crowdsourcing complex work by structuring crowds as organizations. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 3523–3537.
- Van Kleef, G. A., Homan, A. C., & Cheshin, A. (2012). Emotional influence at work: Take it EASI. *Organizational Psychology Review*, 2(4), 311–339. <https://doi.org/10.1177/2041386612454911>
- van Wendel de Joode, R., & Egyedi, T. M. (2005). Handling variety: The tension between adaptability and interoperability of open source software. *Computer Standards & Interfaces*, 28(1), 109–121. <https://doi.org/10.1016/j.csi.2004.12.004>
- Vasilescu, B., Blincoe, K., Xuan, Q., Casalnuovo, C., Damian, D., Devanbu, P., & Filkov, V. (2016). The sky is not the limit: Multitasking across GitHub projects. *Proceedings of the 38th International Conference on Software Engineering*, 994–1005. <https://doi.org/10.1145/2884781.2884875>
- Wageman, R., Hackman, J. R., & Lehman, E. (2005). Team Diagnostic Survey: Development of an Instrument. *The Journal of Applied Behavioral Science*, 41(4), 373–398. <https://doi.org/10.1177/0021886305281984>
- Wegner, D. M. (1987). Transactive Memory: A Contemporary Analysis of the Group Mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of Group Behavior* (pp. 185–208). Springer New York. https://doi.org/10.1007/978-1-4612-4634-3_9
- Wegner, D. M. (1995). A Computer Network Model of Human Transactive Memory. *Social*

- Cognition; New York, 13(3)*, 319–339.
<http://dx.doi.org.proxy.library.cmu.edu/10.1521/soco.1995.13.3.319>
- Weick, K. E. (1995). *Sensemaking in organizations*. Sage Publications.
- Wilden, R., Hohberger, J., Devinney, T. M., & Lumineau, F. (2019). 60 Years of March and Simon's Organizations: An Empirical Examination of its Impact and Influence on Subsequent Research. *Journal of Management Studies, 56(8)*, 1570–1604.
<https://doi.org/10.1111/joms.12531>
- Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a Collective Intelligence Factor in the Performance of Human Groups. *Science, 330(6004)*, 686–688.
- Xie, L. (2021). *A beginner's guide to DAOs*.
https://linda.mirror.xyz/Vh8K4leCGEO06_qSGx-vS5lvgUqhqkCz9ut81WwCP2o
- Zucker, R. (2019). Why Highly Efficient Leaders Fail. *Harvard Business Review*.
<https://hbr.org/2019/02/why-highly-efficient-leaders-fail>