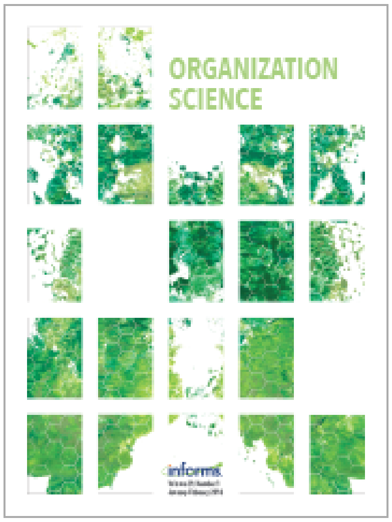

Organization Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

The Dual Challenge of Search and Coordination for Organizational Adaptation: How Structures of Influence Matter

Özgecan Koçak, Daniel A. Levinthal, Phanish Puranam

informs
https://pubsonline.informs.org/journal/orsc

ORGANIZATION SCIENCE
Vol. 34, No. 2, March–April 2023, pp. 851–869
ISSN 1047-7039 (print), ISSN 1526-5455 (online)

The Dual Challenge of Search and Coordination for Organizational Adaptation: How Structures of Influence Matter

Özgecan Koçak,[a,*] Daniel A. Levinthal,[b] Phanish Puranam[c]

[a] Goizueta Business School, Emory University, Atlanta, Georgia 30322; [b] Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104; [c] Strategy, INSEAD, Singapore 138676, Singapore
*Corresponding author

Contact: ozgecan.kocak@emory.edu, https://orcid.org/0000-0002-6974-2382 (OK); levinthal@wharton.upenn.edu, https://orcid.org/0000-0002-8740-6091 (DAL); phanish.puranam@insead.edu, https://orcid.org/0000-0002-0032-8538 (PP)

Received: September 12, 2020
Revised: July 31, 2021; February 10, 2022
Accepted: February 28, 2022
Published Online in Articles in Advance: August 2, 2022

https://doi.org/10.1287/orsc.2022.1601

Abstract. Organizations increasingly need to adapt to challenges in which search and coordination cannot be decoupled. In response, many have experimented with "agile" and "flat" designs that dismantle traditional forms of hierarchy to harness the distributed knowledge of specialized individuals. Despite the popularity of such practices, there is considerable variation in their implementation as well as conceptual ambiguity about the underlying premise. Does effective rapid experimentation necessarily imply the repudiation of hierarchical structures of influence? We use computational models of multiagent reinforcement learning to study the effectiveness of coordinated search in groups that vary in how they influence each other's beliefs. We compare the behavior of flat and hierarchical teams with a baseline structure without any influence on beliefs (a "crowd") when all three are placed in the same task environments. We find that influence on beliefs—whether it is hierarchical or not—makes it less likely that agents stabilize prematurely around their own experiences. However, flat teams can engage in excessive exploration, finding it difficult to converge on good alternatives, whereas hierarchical influence on beliefs reduces simultaneous uncoordinated exploration, introducing a degree of rapid exploitation. As a result, teams that need to achieve agility (i.e., rapid satisfactory results) in environments that require coordinated search may benefit from a hierarchical structure of influence—even when the apex actor has no superior knowledge, foresight, or capacity to control subordinates' actions.

Funding: Puranam thanks the Desmairis Fund at INSEAD for supporting the Organizations and Algorithms project.
Supplemental Material: The online appendix is available at https://doi.org/10.1287/orsc.2022.1601.

Keywords: organizational learning • organization design • computer simulations • computational experiments

1. Introduction

The popularity of concepts and practices, such as the lean start-up and agile teams, indicates that we live in an era that prizes iterative experimentation and real-time adaptation over planning and control (Blank 2013, Denning 2018). The motivations for this shift seem inescapable, such as the rapidity of changes in business challenges that firms face as well as the increase in the richness of feedback and reduced cost of experimentation that contemporary technologies (such as beta testing of novel software or rapid prototyping with three-dimensional printing) provide. Whether firms can succeed in organizing to meet the challenges and opportunities of this era, however, is not clear. As iterative experimentation can rarely be carried out by unitary actors, in many contexts, the push for real-time adaptability presents an organizational (i.e., a multiagent) design challenge.

This challenge is formidable because interdependence between agents as they learn requires organizations to contemporaneously solve the dual challenges of search and coordination. As individuals search for favorable alternatives, the feedback they get depends on the simultaneous search behavior of other agents. For instance, as members of a product design team search for the best course of action to pursue for their individual modules, the performance of their module as well as the overall system and therefore, what they learn from their search will depend not only on their individual search activity but also, on the search activities of other team members. The process of learning for each agent is essentially coupled to that of other agents. In such situations of coupled learning, search and coordination go hand in hand (e.g., Knudsen and Srikanth 2014, Puranam and Swamy 2016, Aggarwal et al. 2017).

In situations when the problem of search is not trivial, no individual agent can lead through superior wisdom. Self-organization through simple imitation of successful actions is also precluded to the extent that specialization limits mutual observability and interpretability of actions (Arrow 1974, Mintzberg 1979). Nonetheless, the ability of individuals in organizations to learn from the beliefs of others may create a possible path to coordination. Organizational agents can exchange beliefs (intentions for actions) and coordinate their actions through convergence of beliefs. For instance, members of a product design team can exchange views on how the various components of the product should relate to each other and what kinds of design choices of each component might optimize overall system performance. Sharing of beliefs enables agents to learn both from their own search experiences and from the search of other agents, thus coordinating their subsequent search activity through converging beliefs.

This raises the question of how exchange of beliefs should be organized. If no agent can have superior wisdom ex ante, the problem is one of creating a structure within which agents that are equally (un-)informed can learn together. The structure of belief exchange and influence can vary from very egalitarian forms where team members have equal influence on one another to one in which a team leader imposes his or her own beliefs on team members, corresponding to what Simon (1947; 1981) referred to as control through decision premises. Which of these might work better in situations of coupled learning? The answer to this question is by no means theoretically obvious. Although interdependence and the need for coordination would suggest that hierarchical structures (with their capacity for producing convergence) should be useful, the lack of superior insight at the apex of the hierarchy (Keum and See 2017) can amplify the risk of convergence to a suboptimal solution (e.g., Knudsen and Srikanth 2014). In fact, classical theorists, such as Thompson (1967, p. 133), have been skeptical about the ability of hierarchical superiors to effectively coordinate actions in situations of complex, unknown, or changing interdependence—the very conditions under which coordinated search is relevant.

Although there is notable prior art on how structures of belief exchange and social influence contribute to organizational adaptation, these do not provide solutions for organizations facing the challenge of simultaneous search and coordination. For instance, there is a large and vibrant literature on how solutions to search problems diffuse within a network of actors (e.g., March 1991, Lazer and Friedman 2007, Fang et al. 2010, Mason and Watts 2012, Shore et al. 2015, Brackbill and Centola 2020). However, the payoff or measures of performance in these models entail no interdependencies between the actions of the various actors. There is also a sizeable body of theoretical and empirical work on how social collectives converge on shared action and language as beliefs and norms diffuse through networks of influence (e.g., Friedkin and Johnsen 1990, Friedkin 2006, Centola and Macy 2007, Mason et al. 2007, Centola 2015, Guilbeault et al. 2018). However, these pure coordination tasks do not involve a search for inherently superior solutions. Prior studies of the combined challenge of search and coordination have assumed these to be separable across agents or in time (e.g., Rivkin and Siggelkow 2003). The few papers that have examined the problem of simultaneous search and coordination (e.g., Lounamaa and March 1987, Knudsen and Srikanth 2014, Puranam and Swamy 2016) have done so mostly in dyadic systems.

As a consequence, which structures of influence help to accomplish coordinated search remain unknown. Put differently, we know that a fundamental role of organizations is to facilitate coordination (Thompson 1967, Weick 1974, Nelson and Winter 1982, Puranam 2018); yet, how organizations accomplish coordination when search is simultaneously important is an open question. This gap in our understanding also has practical significance. Consider current interest in agile development practices, which originated in the software industry but are now being applied far afield (e.g., Hill 2020). These involve an iterative approach of using frequent meetings to prioritize among multiple goals (Wu and Ghosh 2021), representing an explicit recognition by practitioners that search and coordination cannot be profitably separated in collective experimentation (Rigby et al. 2016, Denning 2018, Schwaber and Sutherland 2020).

Although agile organizing has been popular among practitioners and has also received attention from scholars, there is considerable variation in beliefs about what the structure of influence within agile teams should entail and how these structures should be implemented. Some accounts stress the importance of densely connected egalitarian networks with no hierarchy (Denning 2018) and suggest that the introduction of hierarchical elements is one of the reasons that agile practices may fail (Rigby et al. 2016). However, the Scrum Guide, which defines the minimal standards for the most popular method of agile organizing, explicitly recommends roles that might exert asymmetric influence (Schwaber and Sutherland 2020). Others emphasize the role that "scrum masters" and "product owners" play in achieving coordinated action in agile teams (Dönmez et al. 2016) or the need for particular forms of hierarchical influence (e.g., servant leaders) for agile development to be effective (Guinan et al. 2019, Hill 2020).

Using a computational agent-based model, we explore some of the key underlying relationships between structures of influence and organizational adaptation. In particular, we ask: when do hierarchical structures of influence aid organizational adaptation?

The prescriptive writing on agile teams that advocates a dense and symmetric influence structure (henceforth, a flat team; e.g., figure 1-1 in Denning 2018) is consistent with the skeptical view on the value of hierarchy for organizational adaptation. The implicit claim in such writing is that lateral social influence in flat teams suffices for coordinated search without any need for hierarchical direction. However, as we have noted, other advocates of agile practices recognize that team leaders might exist, with greater influence on all other members. To understand the underlying mechanisms and contingencies that explain which structures of influence might aid organizational adaptation, we contrast flat with hierarchical teams.

We note that hierarchical in our usage refers only to influence in our models stripped of other critical features associated with administrative hierarchies, such as fiat (Williamson 1975), incentive schemes defined by a principal (Jensen and Meckling 1976), and a vertical division of decision rights and labor (Galbraith 1973) among others. Instead, we focus on the idea that even when leaders do not have the wisdom or the ability to effectively control their subordinates' actions, they may lead by influencing their beliefs. As Simon (1981, p. 143) noted, influencing beliefs is a key aspect of administrative hierarchy and is a feasible way to exercise authority even in knowledge-intensive contexts where a boss cannot directly control actions.

As a benchmark to both hierarchical and flat teams, we also consider a limit case in which all agents act without any belief exchange among themselves—as a crowd. Unlike in the wisdom of crowds literature that focuses on one-shot aggregation of estimates (Surowiecki 2005, Page 2010), here we use the crowd as a useful benchmark against which to illustrate the impact of social influence in a dynamic learning task that involves both search and coordination. Further, by varying the influence weights that each agent places on updating their beliefs based on their own learning experience and the social influence of connected others, we are able to smoothly cover the space of influence structures bounded by these three archetypes—flat teams, crowds, and hierarchical teams.

We find that exploration is facilitated by social influence. When agents are influenced by the distinct beliefs of other agents during the search process, their own beliefs are less likely to stabilize prematurely around their own experiences, which in turn, introduces more variety in the explored alternatives. This property is critical to avoiding fixation on inferior outcomes, and both flat and hierarchical teams produce this benefit relative to crowds. However, flat teams can engage in excessive exploration, finding it difficult to settle down even on a previously sampled "good" set of interdependent actions. When social influence is hierarchical, it reduces simultaneous uncoordinated exploration as influential actors vary their actions less often than the influenced actors. As a result, hierarchical influence introduces a degree of exploitation—or reduced variety in the collective alternatives selected. This property gives hierarchical teams an advantage when the benefits of rapid convergence to good, if not optimal, choices outweigh the costs of convergence to a nonglobal peak. Thus, although the importance of exploration and exploitation in adaptation processes is well understood (Holland 1975, March 1991, Miller and Page 2009), the subtle mechanisms through which hierarchical influence structures produce this balance in multiagent learning systems that face both search and coordination challenges are unique to our analysis.

An implication is that teams that need to achieve agility—in the sense of rapid results—in environments that require coordinated search benefit from a hierarchical structure of influence. Our results thus offer a perspective on why hierarchical influence structures—whether realized through classical authority relationships (March and Simon 1958); more recently popularized "soft authority," such as coaching roles and servant leaders (Laloux 2014, Lee and Edmondson 2017); or informal status hierarchies (Bunderson et al. 2016, Greer et al. 2018)—may persist even when the apex agent has no superior knowledge, foresight, or authority as long as organizational adaptation unfolds in task environments in which both search and coordination are important.

2. Searching in Contexts: Environments and Structures of Influence

We first discuss prior modeling work on the different ways in which the environment within which the organization adapts has been characterized and then turn to the relationship between different structures of influence within an organization and its ability to adapt.

2.1. Search and Coordination as Challenges to Organizational Adaptation

In a world with high interdependence—of imperfect decomposability of task structures (Simon 1962)—an important aspect of organizational adaptation is the search by agents for superior collective actions (i.e., a set of mutually compatible actions across agents that in aggregate improve organizational performance) (Nelson and Winter 1982, Henderson and Clark 1990). The literature on fitness landscapes (Levinthal 1997) has highlighted the implications of interdependence between actions for the *ruggedness* of the payoff landscape—the presence of multiple peaks. However, in principle, the variability of payoffs and interdependence between the actions of different actors are independent properties. Situations with high interdependence between actions can coexist with no variability (e.g., all peaks are equivalent, but what matters is that actors converge to the same peak). At the other extreme, peaks can vary in value, and the returns to each actor for finding a good peak are independent of the choices of others. Finally and one would expect more typically, there are situations where interdependence between agents also matters, and the points of coordination vary in their payoff.

Although some organizational models have recognized that search and coordination are both important to adaptation, they have assumed that these two aspects of adaptation can be decomposed and performed at different time periods and/or by different agents. In this spirit, Rivkin and Siggelkow (2003) and Siggelkow and Rivkin (2005) examine two-layered hierarchies of a superior and two subordinates, each of whom searches a different part of the same rugged landscape. Superiors have the ultimate decision rights to accept or reject proposals from subordinates but do not influence their beliefs. The assumption here is that search and coordination can be separated in a two-stage process, an assumption that is challenged in many contexts (Rigby et al. 2016, Denning 2018).

The search for superior interdependent actions has been modeled in dyadic models of coupled learning. In such models, each agent, possibly specializing in a particular domain, is able to search for superior alternatives. The payoffs to their choices, however, depend on what other agents simultaneously choose (Lave and March 1975, Lounamaa and March 1987, Puranam and Swamy 2016, Aggarwal et al. 2017). Of particular interest from the perspective of the current work are models of coupled search by pairs of agents whose mental models of the search space can be shaped both by feedback from the environment and by mutual influence through communication with the other agent. Knudsen and Srikanth (2014) show that communication helps to overcome mutual confusion by aligning agents' beliefs about the search landscape, but communication beyond moderate frequency causes joint myopia as agents prematurely converge on a narrow portion of the search space. Puranam and Swamy (2016) demonstrate the benefits of common priors in situations involving both search and coordination as a safeguard against superstitious learning (March and Levinthal 1993). These models do not, however, compare influence structures or systems larger than a dyad.

2.2. Structures of Influence Within Organizations

Influence refers to a change in beliefs and consequently, behavior of one agent because of the actions (including acts of communication) of another. We use the term influence in this paper always as influence on beliefs to distinguish from interdependence between actions. The strength and direction of influence relations within organizations can vary significantly from situations where agents are symmetric sources of information to each other (e.g., an advisory or friendship relation) all the way to settings in which one actor has complete control of actions of another (i.e., if the beliefs of a hierarchical superior completely determine the beliefs of a subordinate, effectively creating a situation of fiat).

The work of March (1991) on organizational exploration/exploitation featured this property and has served as an important touchstone for multiagent models of organizational adaptation and learning featuring influence relationships among agents. In this model, agents adapt their beliefs to those of others in the organization, whereas the organization also learns from the higher-performing individuals. The influence structure is modeled in a reduced form through an organizational code. Individual agents learn from the code, and the code learns from high-performing agents, although at different rates. March (1991) shows that, in this model, organizational knowledge is highest when the organizational code updates relatively fast and individuals learn relatively slowly from the code.

This implicit hub-and-spoke (code and agents) topology was extended in later work to explicitly model connections between agents. Miller et al. (2006) find that the level of organizational knowledge attained is superior when the influence structure is equivalent to a network with moderate degrees of clustering. Fang et al. (2010, p. 628) state that "semi-isolated clusters enable the creation and preservation of heterogeneous ideas, whereas a few external links help the best ideas to diffuse." Lazer and Friedman (2007) and Schilling and Fang (2014) explicitly model influence structures as networks and compare basic archetypes of symmetric tie networks (e.g., linear, random graph, "caveman") with one another. Overall, these studies highlight a tension close to the spirit of the exploration/exploitation trade-off: densely connected topologies do poorly at maintaining the diversity required for finding high-performing strategies but do well at rapidly disseminating them through diffusion.

Influence structures in organizations often have hierarchical patterns—denoting asymmetric, transitive, acyclic relationships (Simon 1962, Ahl and Allen 1996). For instance, in a formal administrative or informal status hierarchy, A has influence over B, who has influence over C (and so, A indirectly has influence over C).

Admittedly, influence is but one aspect of an administrative hierarchy, although arguably a crucial one.[1] Indeed, central to the view of administrative behavior of Simon (1947) is the role of decision premises influencing lower-level actors' beliefs and behavior. Extending this assumption to multiple agents, their actions can effectively be coordinated without the need for them to communicate directly with each other. Simon (1981) elaborated on this idea to note that when uncertainty is likely to affect many parts of an organization in the same way, "it may be advantageous to centralize the making of assumptions about the future and to require the decentralized units to use these assumptions in their decisions" (Simon 1981, p. 43).

It is intuitive that if the apex actors in a hierarchical structure of influence are particularly knowledgeable, their downward influence can benefit the organization. More generally, the direction and strength of influence in many prior models are often assumed to be correlated with agents' competence at search or the quality of information they have (March 1991, Miller et al. 2006, Lazer and Friedman 2007, Schilling and Fang 2014). A similar link is present in work in economics (Bala and Goyal 1998) and swarm models in biology (Kennedy and Mendes 2002, List et al. 2009) that incorporate success-biased social influence. Experimental studies of parallel search in networks with human subjects also explicitly enable agents to imitate others that are successful (Mason and Watts 2012, Acerbi et al. 2016). This coupling between influence and competence also exists in models in which centralized authority is used to resolve conflicts (Rivkin and Siggelkow 2003), coordinate (Mihm et al. 2010), or solve rare and difficult problems (Garicano 2000) that an organization encounters as it adapts. In fact, social psychologists have argued that correlation of expertise and rank in hierarchy may be necessary for the functionality of hierarchical structures (Anderson and Willer 2014). For instance, Tarakci et al. (2016) find in their computational modeling as well as a laboratory study and a field study, that intragroup differences in power and influence improve performance when they are correlated with task competence and harm group performance when they are held constant or are not aligned with task competence.

One possibility is to read these prior studies as implicitly ruling out any possible advantage for hierarchical influence for organizational adaptation unless this influence structure is correlated with expertise. This could explain the abhorrence toward hierarchy in agile development practices, where the joint importance of search with coordination makes any superior wisdom at the apex unlikely. However, another possibility is to attribute earlier findings to the particular task environment that was considered. Many of the studies of organizational adaptation noted feature parallel search, in which the outcome of an agent's actions does not depend on the actions taken by others, even though they may learn from each other (March 1991, Lazer and Friedman 2007, Fang et al. 2010, Mason and Watts 2012, Schilling and Fang 2014, Tarakci et al. 2016). Although the search by each agent may take place in a task environment where payoffs vary, these payoffs are not driven by any interdependence between the actors (although interdependence may exist between the multiple actions taken by a single actor) (e.g., Lazer and Friedman 2007). In contrast, we examine task environments in which the problem of interdependent action or coordination arises in varying degrees, in conjunction with the process of search.

Although the models of search noted connect influence with expertise, prominent models of influence in sociology do not purport to deal with adaptation to an environment and therefore, leave expertise undefined. Instead, they have been concerned with pure coordination (that is, whether social influence processes lead to convergence on a common belief or not depending on the social structures of influence and the motives of interacting agents) (e.g., Centola and Macy 2007, Mason et al. 2007, Centola 2015, Centola and Baronchelli 2015, DellaPosta et al. 2015, Guilbeault et al. 2018). Computational and empirical studies of asymmetric influence relationships in these contexts find that asymmetry facilitates coordination (e.g., Friedkin and Johnsen 1990, Friedkin 2006).

The existing literature thus indicates that hierarchical structures of influence may have advantages either when coordination (not search) is the key imperative or in search situations when the more influential agents have superior knowledge or capabilities. However, whether hierarchical influence is an advantage or a liability when both search and coordination matter—so that hierarchical superiors cannot be assumed to have more accurate beliefs about the task environment—remains largely unexamined.

2.3. Summary: Generalizing Task Environments and Influence Structures

To improve our understanding about how different structures of influence aid organizational adaptation in different task environments, we require two ingredients. First, we need a framework to model different structures of influence on beliefs in multiagent systems, and for this, we draw on network-based representations, as prior scholars have done (e.g., Friedkin and Johnsen 2002, Lazer and Friedman 2007, Fang et al. 2010).

Second, we need to represent varying contexts of organizational adaptation (see Table 1). One fundamental challenge for organizational adaptation is posed by the degree to which alternatives vary in their payoffs.

Table 1. Organizational Task Environments: Payoff Variability and Interdependence

		Interdependence	
		No interdependence between agents	Interdependence between agents
Payoff Variability	Landscapes with variable payoffs: High penalty for converging to local peak	Pure search challenge e.g., parallel problem solving and innovation contests (Lazer and Friedman 2007; Fang et al. 2010)	Challenge of search with coordination e.g., formation of routines, adaptation to architectural change (Puranam, Raveendran and Knudsen 2012; Knudsen and Srikanth 2014)
	Uniform landscapes: Low penalty for converging to local peak	*Neither search nor coordination pose challenges to organizational adaptation*	Pure coordination challenge e.g., formation of conventions and codes of communication (Friedkin and Johnsen 2002; Spike et al. 2017)

We term this dimension as variability, and it tunes the challenge of search. The second basic dimension on which environments might differ is the degree to which agents' payoffs depend on the actions of others (i.e., interdependence). As indicated, the effect of hierarchical influence on beliefs has previously been studied in environments that alter the variability of payoffs in the absence of interdependence between agents (e.g., Lazer and Friedman 2007, Fang et al. 2010) or in the context of social influence in a low-variability, pure coordination problem (e.g., Friedkin and Johnsen 2002, Centola and Baronchelli 2015) but not when the two problems coexist. Arguably, however, this conjunction of these dual challenges is central to most common adaptation problems in organizations given the reality of division of labor and resulting interdependence in actions across agents (Thompson 1967, Puranam et al. 2012).

We next describe how these two elements are captured in our model.

3. Model Structure

We conceive of an organization as an adaptive system comprising multiple individuals who are themselves adapting to their individual but possibly, interdependent task environments. Each actor chooses one of many possible actions, and a set of individual actions collectively represents an *organizational action*. The central problem we analyze is how this adaptive process is affected by the structure of influence on beliefs among agents and properties of the task environment (i.e., the variability of payoffs and interdependence in payoffs across actors). We abstract away from well-understood agency problems (i.e., trade-offs between incentives and risk bearing and between group and individual incentives) by assuming that each agent's payoff is positively correlated with the organizational payoff. We focus on the significant organizational challenges posed by search and coordination even when there are no challenges to cooperation—on the adaptive process through which superior organizational actions are discovered collectively (e.g., March 1991, Levinthal 1997, Ethiraj and Levinthal 2004, Clement and Puranam 2018).[2]

The description covers the baseline model from which we derive our results. In additional analysis, we examine several alternative formulations of the model and the robustness of our results.

3.1. Task Environment

We set up a task environment that allows us to separately tune the challenge posed by the dual problems of search and coordination. In order to model interdependence between actors in a search space, we have separate parameters to tune the variety of payoffs to alternative choices and the degree to which agents are rewarded for converging on compatible actions (i.e., for coordinating).

We model the challenge of searching for individual actions that contribute to superior organizational actions by changing the distribution of payoffs over actions available to each agent. More formally, each agent faces S discrete choice alternatives. The baseline value (v) of each choice is treated as deterministic and constant over time. This value is the same for any agent who makes that choice. Organizational payoffs are treated as the average of individual payoffs because we assume incentive alignment in this model.

In seeding the payoffs, one choice is randomly picked to be the global peak, which we set at a value of one. The values of the remaining $S - 1$ choices are drawn from a uniform random distribution with a range of $[\delta_0, \delta_1]$, where δ_0 lies in the interval [0, 1] and δ_1 is in $[\delta_0, 1]$. This implies that although the payoff of the global peak is always one, the payoff of the average suboptimal action is $(\delta_0 + \delta_1)/2$. At low values of δ_1, the global peak will be significantly better than the next best alternative, although at a high value of δ_1, the peak is hard to distinguish from other choices. For instance, a big difference between the global peak and the next best peak (a low δ_1) corresponds to a winner takes all market, where low search and switching costs create a very uneven competitive outcome for producers that can find the peak and those that cannot. Conversely, a high δ_1 corresponds to markets where high differentiation and high switching costs create many viable peaks for firms to climb. In contrast, a high value of δ_0 creates a task environment where even the worst outcome is not so much worse than the best outcome. This is a task environment where it is difficult to fail, as in the case of teams tasked with making incremental innovations in a firm that already has a captive audience. Trapping to suboptimal actions is, therefore, more harmful when $(\delta_0 + \delta_1)/2$ is smaller.

To tune the need for coordination between agents' actions requires a specification of which set of agent's actions is to be considered as being mutually compatible and the payoff advantages to finding compatible actions. We assume that agents' actions have converged (or, equivalently, are coordinated) if they select actions with the same index (e.g., agents 1 and 2 both select action 3 of S possibilities). This should be interpreted as taking mutually compatible actions, not necessarily identical actions: for instance, when they adopt a complementary division of labor or find interdependent patterns of action in the case of organizational routines.[3] We model the importance of coordination as a weight on the baseline value of the individual actor's choice and the measure for the degree of convergence among the set of actors in their choices. Formally, the *payoff* (π) an agent i receives when picking choice k at time t is given by the equation

$$\pi_{i,k,t} = v_k[1 - \chi(1 - F_{k,t})], \quad (1)$$

where $F_{k,t}$ = the extent of convergence (the fraction of other agents who pick the same choice k at time t), v_k = baseline value of choice k, and χ = convergence weight in the range [0, 1].

In effect, we create a performance penalty for the organization (and therefore, each individual) when agents take incompatible actions. In an environment in which coordination does not matter ($\chi = 0$), the contribution of each individual's action to organizational payoff does not depend on the choices of the other agents but only on that individual's own baseline payoffs v_k. Increasing χ increases the penalty to nonconvergence of actions, which thus tunes the importance of coordination. On the other hand, when $F_{k,t} = 1$, $\pi_{i,k,t} = v_k$ for any χ. In Online Appendix Section 1, we give a detailed illustration of how χ tunes the payoff landscape in a two-choice world.

Seen as complex adaptive systems, organizations adapt when their components adapt to the task environment and each other. Pure search processes involve no mutual adaptation among subunits, and pure coordination processes involve no adaptation to the task environment. The more general form of this problem is when both are necessary. Each of these cases can be captured in our task environment. When $\delta_0 < \delta_1 \leq 1$ and $\chi = 0$, the key challenge for organizational adaptation is variability in payoffs, not interdependence among agents—this is a case of parallel search (the upper left corner in Table 1). This would be a setting akin to a bandit model with multiple agents receiving feedback from the task environment independently of each other. At the polar opposite is the case when $\chi = 1$ and $\delta_0 = \delta_1 = 1$, where the challenge of coordination of actions is relevant but which organizational action is coordinated on does not matter (the lower right corner in Table 1). This is akin to a pure matching game (Schelling 1960). In between these extremes is the broad range of parameter values in which the challenges of search and coordination are both relevant (i.e., $\chi > 0$, $\delta_0 < \delta_1 \leq 1$) (the top right corner in Table 1).

3.2. Agents' Beliefs

Each agent i maintains a time-varying vector of beliefs $\boldsymbol{b}_{i,t}$ of dimension S whose elements $b_{i,k,t}$ in [0, 1] denote how attractive a particular choice, k, is at time t for individual i. The beliefs can be thought of as mental representations regarding the relative desirability or "attractiveness" of the alternative actions (Sutton and Barto 1998, Puranam and Håkonsson 2015). The initial vector of these beliefs, $\boldsymbol{b}_{i,0,}$ is formed by taking a draw from a uniform random distribution with range [0, 1] for each possible action. An agent's $\boldsymbol{b}$ vector changes over time based on the payoffs received as well as through the influence of beliefs held by other agents.

First, each agent's payoff determined through (1) is used to update their individual assessment of attractiveness $b'_{i,k,t}$, of the choice k they made at time t for that agent, by averaging the payoffs received with previous beliefs (Sutton and Barto 1998, Posen and Levinthal 2012):

$$b'_{ikt} = (b_{ikt-1} \times (n_{ijt} - 1) + \pi_{ikt})/n_{ijt}, \quad (2)$$

where n_{ijt} is the number of times that i has sampled k by period t. In additional analysis, we also consider an update process based on a tunable formulation of reinforcement learning (Bush and Mosteller 1955), which also features the behavioral property known as the law of recency, whereby recent feedback is weighed more heavily than past feedback (Erev and Roth 1998, Puranam and Håkonsson 2015).

Second, each individual's final belief vector $\boldsymbol{b}_{i,t}$ is formed on the basis of both their individual assessments $(\boldsymbol{b}'_{i,t})$ as well as the average of the individual assessments of all other agents who have an influence on the focal agent (Friedkin and Johnsen 1990, Becker et al. 2017, Friedkin et al. 2019). This updating process is standard in models of social influence and originates in the work of De Groot (1974). The final (postinfluence) belief vector for the focal agent i at time t is specified as

$$\boldsymbol{b}_{\mathrm{i,t}} = w \times \boldsymbol{b}'_{\mathrm{i,t}} + (1 - w) \times \boldsymbol{b}'_{\mathrm{j,t \in Ni}}, \tag{3}$$

where w, with range [0, 1], indexes the weight that each agent places on their own *individual assessments* when updating their beliefs ($\boldsymbol{b}'_{\mathrm{i,t}}$) and $\boldsymbol{b}'_{\mathrm{j,t \in Ni}}$ indicates the average of the individual assessment vectors of agent i's N_{i} network neighbors at time t.

DeGroot updating is likely to approximate social influence in settings where we can expect individuals to communicate about the beliefs (decision premises) that underlie their actions—a key and distinguishing feature of influence processes within organizations (Argote 2012; Simon 1947, 1981).[4] Despite its apparent simplicity, (3) captures several important aspects and consequences of belief sharing. First, by tuning the influence weight, the extent of relative influence can be varied rather than discretized through setting the direction of a tie. Second, this rule accounts for a form of social reinforcement; when agents' beliefs on a specific alternative are aligned (e.g., both agents believe the alternative j is attractive), those beliefs are likely to be strengthened compared with misaligned beliefs. Third, this rule can also produce creative outcomes through recombination of beliefs—an option that was not the most attractive for either agent can become so after belief sharing.

Choice over the resultant belief vector $\boldsymbol{b}_{\mathrm{i,t}}$ follows the standard softmax action selection rule with an exploration parameter τ (Luce 1959, Sutton and Barto 1998), where the probability of an agent (we suppress agent index) choosing alternative k at time t is

$$p_t(k) = \frac{e^{b_{k,t}/\tau}}{\sum_{k=1}^{S} e^{b_{k,t}/\tau}}. \tag{4}$$

As the parameter τ approaches zero, even small differences in beliefs lead to very divergent probabilities of the choice of action, and in that sense, low values of τ represent an individually exploitative search strategy. In contrast, as τ takes on larger values, the choice of behavior becomes less sensitive to differences in beliefs and is closer to random (i.e., high individual exploration).

3.3. Interdependence and Influence Structures

In task environments where agents' payoffs depend on other agents' actions, the shape of the interdependence structure can impact the evolution of beliefs. In baseline models, we assume a global interdependence structure, whereby each agent's contribution to the group outcome is equal and each agent's payoff depends on the group outcome in equal measure. In supplemental models, we implement more nuanced structures where the structure of interdependence follows the structure of influence.

Although interdependence structures moderate learning from the task environment, learning from other agents is determined by the influence structure. We model three distinct structures of influence (Figure 1) by varying w in Equation (3) in a fully connected structure of influence. As a baseline, crowds are sets of agents that lack any social influence ($w = 1$). They represent the average behavior of isolated agents who might nonetheless be in a task environment that requires coordination.[5] Flat teams are fully connected symmetric networks, where w is the same for every agent (0.5 in baseline models, meaning that agents weigh their own priors and the connected agents' beliefs equally). Hierarchical teams are like flat teams except that one agent—the "apex agent" or "leader"—has higher w than the rest, meaning that it has asymmetric influence on each of the other team members. In baseline models, we set the apex agent's $w = 1$, creating an acyclic structure of asymmetric influence, meaning that no other agent can influence this apex agent directly or through any other agent.

Dyadic asymmetries in influence are abundant in organizations and can arise from differences in power, prestige, status, network position, and of course, formal authority. Acyclicity, however, does not always arise spontaneously and is most often a formal design feature in organizations (Chandler 1990, Krackhardt

Figure 1. Alternative Influence Structures

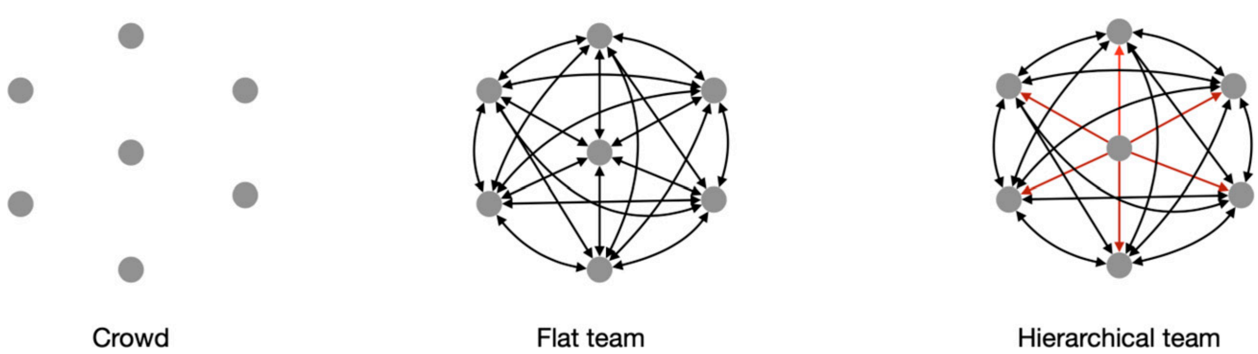

1994, Martin 2009). In supplemental models, we explore intermediate forms besides the three archetypes by varying the number and placement of dyadic ties with asymmetric influence (thus, exploring the effects of dyadic asymmetry and acyclicity), the apex agent's w in the hierarchical team (thus, exploring the effect of varying *degrees* of hierarchical influence), and w for all agents (thus, exploring the range of structures between crowds and flat teams with varying strengths of social influence relative to individual learning from the task environment).

4. Baseline Results

The fundamental attributes that may impact the efficacy of the alternative influence structures are attributes of the task environment and properties of individual learning. With regard to the former, the critical parameters are interdependence as determined by χ and variability of payoffs as characterized by δ_0 and δ_1. In terms of attributes of individual agents, the key parameter is the individual balance of exploration/exploitation, τ. In baseline models, the agent exploration parameter is set at a moderate level of $\tau = 0.05$ (see Table 2 for a summary of these four key parameters and their ranges).

We examine the three archetypical influence structures in the basic analysis holding the number of actors, N, to seven. We also set $S = N$ to allow for maximum possible diversity of actions to begin with (i.e., each agent could in the limit prefer a different action initially; in the analysis, their initial preference is randomized). For each influence structure, we present results averaged across 1,000 randomly sampled task environments.

Figure 2 provides a comparison of relative cumulative performance of the three archetypes within the basic task environments described in Table 1. The graphs show the parameter space for all values of δ_0 [0, 1] and all values of δ_1 [δ_0, 1] in steps of 0.1. The constraint that δ_1 lies in [δ_0, 1] gives the graphs their triangular shape. We show results for three values of interdependence at $\chi = \{0, 0.5, 1\}$ across the three columns in Figure 2. The color of each cell shows the dominant influence structure: flat team (green), hierarchical team (red), or crowd (blue). The depth of the shading increases with the magnitude of the advantage of the dominant form over the next best form in terms of cumulative performance at time T. The results examine three time slices at $T = \{50, 100, 500\}$. The top right corners of graphs in the first column (where $\chi = 0$ and $\delta_0 = \delta_1 = 1$) correspond to uniform landscapes with no interdependence between agents—that is, task environments where neither search nor coordination pose a challenge. These cells are pale in color because all forms attain identical performance in these environments.

Table 2. Key Model Parameters

Symbol	Range	Description
δ_0	[0, 1]	Variability of payoffs increases
δ_1	[δ_0, 1]	with $1 - (\delta_0 + \delta_1)/2$
χ	[0, 1]	Interdependence increases with χ
τ	>0	Explorative tendency increases with τ

The results in Figure 2 yield a clear general pattern. The dominant form in both the short and long terms in situations of pure parallel search (i.e., with no interdependence, $\chi = 0$) is the flat team. However, once we allow for interdependence in the task environment ($\chi > 0$), then in the short term ($T = 50$), the dominant form is the hierarchical team, whereas in the longer term it is the crowd. In other words, if the forces of selection in the task environments with interdependence are somewhat myopic (Levinthal and Posen 2007), hierarchical influence structures facilitate search better than either crowds or flat teams, whereas crowd-like structures prevail in the longer term.

This pattern of results supports some baseline intuitions, such as the value of flat teams for pure parallel search when no agent has superior insight, but also, challenges others. In task environments that are more characteristic of organizational adaptation (i.e., where search and coordination both matter), there are two puzzling results to understand. First, why do hierarchical teams enjoy an advantage over flat teams at all? By construction, we have ruled out any differential wisdom at the apex of the hierarchical team, nor are there any incentive conflicts to be controlled by the "team leader." Second, why does a zero-influence structure, such as the crowd, dominate both flat and hierarchical teams in the long term?

5. Analysis of Mechanisms

To understand why flat teams are beaten by hierarchical teams in the short term and crowds in the long term, we depict the dynamics of the model in Figure 3 for a task environment in which both search and coordination matter: S = 7, $\delta_0 = 0$, $\delta_1 = 1$, $\chi = 1$. We consider the same three influence structures as in the baseline analysis (flat teams, hierarchical teams, and crowd with $n = 7$) and examine six outcomes over 500 periods to reveal the mechanisms producing these results.

First, we compute the *average payoffs* in [0, 1] received by agents in the organization in each period as in Equation (1). Second, we measure the degree of *convergence* of actors' choices as a Herfindahl concentration index (sum of squared shares of agents across choices). A third measure is the *success at search*, measured by the proportion of groups that reached the global peak in the landscape ($\pi_t = 1$) in each period. The fourth and fifth measures capture *similarity* and *resolution* of beliefs. Similarity of beliefs is a precursor to similarity

Figure 2. Dominant Influence Structure Based on Cumulative Performance at T = 50, 100, and 500 (n = S = 7, τ = 0.05)

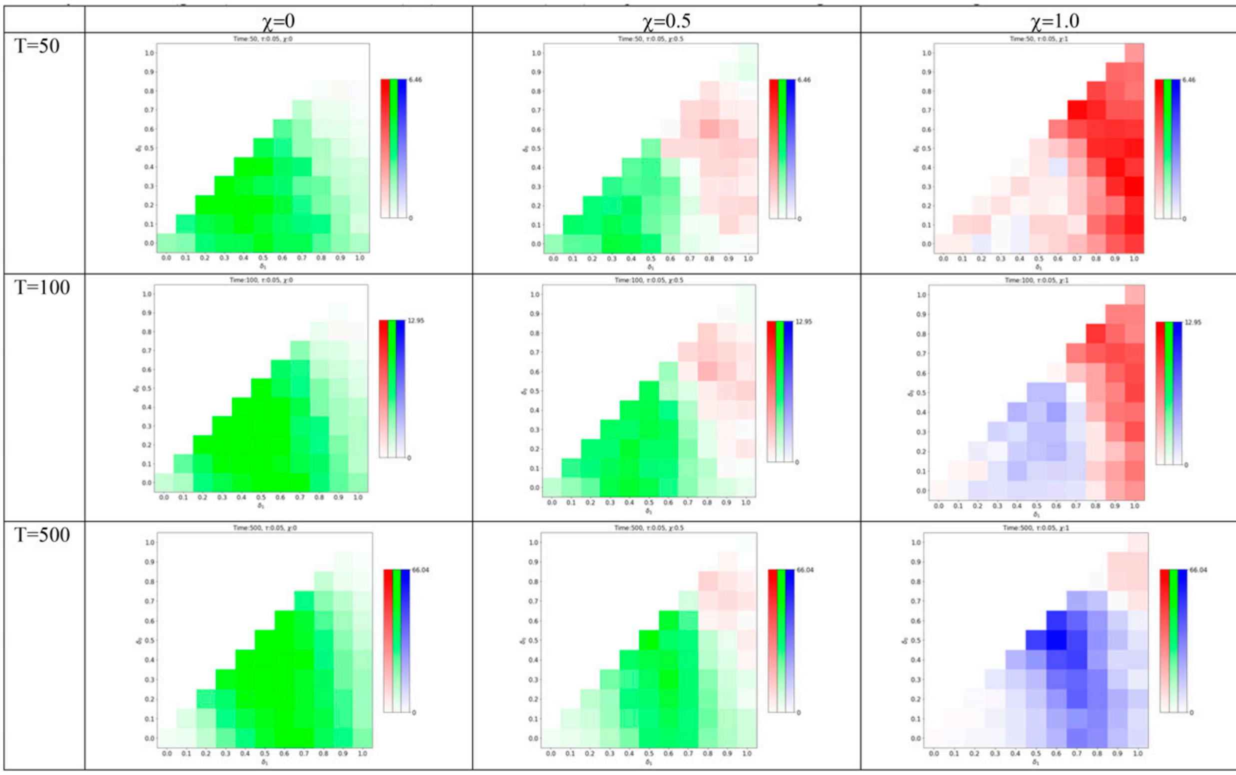

Notes. The dominant structure is indicated by color: flat team (green), hierarchical team (red), or crowd (blue). The depth of shade shows the magnitude of advantage over next best structure.

of actions (or equivalently, the selection of compatible actions). We measure similarity of beliefs across agents as one minus the average of the absolute differences between belief vectors across all pairs of agents. Note that similarity need not be at its maximum value for actions to be identical (see Equation (4)). We measure the resolution (or sharpness) of beliefs as the difference between maximum and minimum belief strengths within an agent's belief vector. Resolution of beliefs shapes the probability with which the action with the most attractive belief is selected (again see Equation (4)). Finally, the sixth measure plots *instability* in behavior. We consider the stability of agents' payoffs and behavior over time (period to period changes) to get a better understanding of the dynamics of the process and in particular, its asymptotic properties (e.g., whether the process converges to a stable fixed point or not).

Figure 3 shows that in this task environment where both search and coordination matter, hierarchical teams produce the highest and most rapid degree of convergence of actions among agents, followed by crowds. Flat teams attain the lowest convergence and display the greatest instability over longer periods of time. This ability to achieve rapid convergence leads to (and is, in turn, reinforced by) an early advantage for hierarchical teams in terms of payoffs. However, there is also a significant probability of such rapidly converging teams getting trapped on the nonpeaks, which can remain more attractive than the peak in the agents' beliefs when a sizable number of agents select it. We can see this in Figure 3 by comparing the graphs for convergence and success at search. Even though hierarchical teams have the highest convergence of actions for the longest time, flat teams and crowds eventually overtake them in terms of success at finding the best peak because hierarchical teams are more likely to get trapped on nonpeaks.[6] The risks of entrapment to a nonpeak are significant when χ increases (see Online Appendix Section 1). This is the reason that in the long run, crowds may overtake hierarchical teams even in task environments with interdependence, although cumulative payoffs can remain higher for hierarchical teams for a very long time. How quickly hierarchical teams are overtaken by crowds in the long run in high-χ environments (third column in Figure 2) depends also on the distribution of alternatives in the landscape (hierarchies retain their advantage longer in munificent task environments).

To highlight the novelty of these results, it is worth noting that when trying to understand the processes that lead to faster and greater convergence of actions

Figure 3. Comparative Dynamics for $\chi = 1.0$, $\delta_0 = 0$, $\delta_1 = 1$, and $\tau = 0.05$ (n = S = 7)

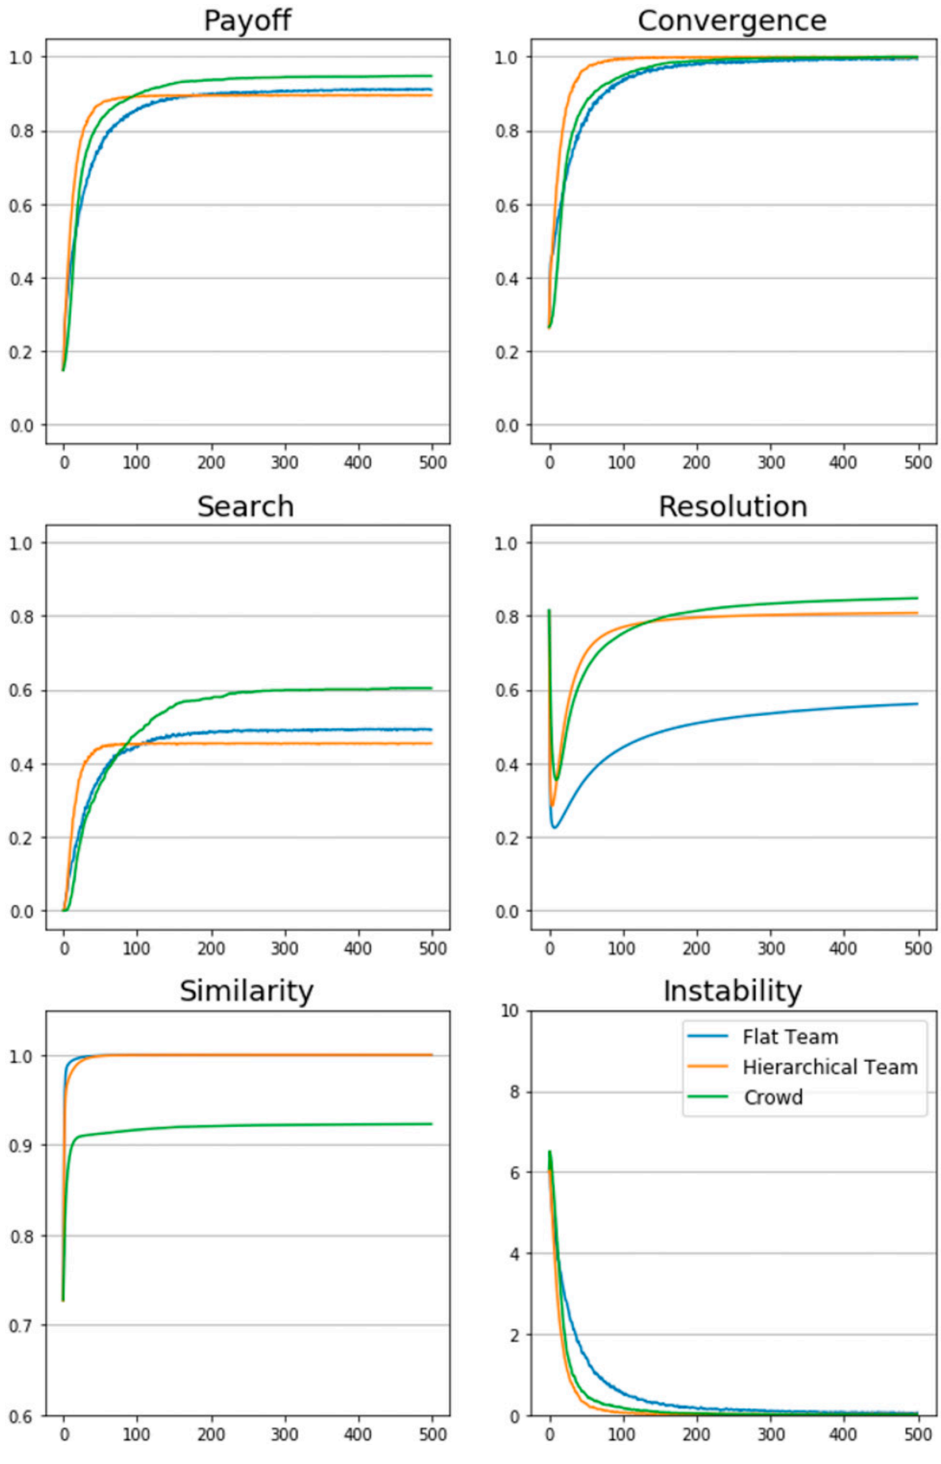

in hierarchical teams, a reasonable prior is that hierarchical social influence increases the similarity of beliefs among agents more rapidly. However, Figure 3 shows that flat teams attain near-perfect similarity of beliefs even faster than hierarchical teams. Why then do hierarchical teams converge faster than nonhierarchical teams? The reason is that similarity of beliefs is not sufficient to produce identical actions. If beliefs are similar but fuzzy (have low resolution), they can lead to diverse actions. Given interdependence, this produces a performance penalty, which further retards convergence. When the task environment features interdependence, both high resolution and similarity are necessary to rapidly increase convergence in actions. When both search and coordination matter, hierarchical influence structures attain rapid convergence because of high-resolution beliefs that are similar across agents. Flat teams accomplish high similarity rapidly, but their members have beliefs with low resolution, whereas agents in crowds develop high-resolution beliefs but do not attain as high levels of similarity. Only hierarchical teams accomplish both.

The persistent low resolution of beliefs in flat teams also explains why crowds outperform flat teams in the long run. Agents in crowds can eventually produce convergent actions by learning from payoffs in high-χ environments. Agents in flat teams, however, overexplore because of their low-resolution beliefs. This is visible in Figure 3 in the panel that shows instability of actions, where we see that members of flat teams are more likely to change their actions from one period to the next compared with agents in hierarchical teams or crowds. In effect, the low resolution of beliefs produced by the flat team acts as a functional equivalent to agents having an individual tendency to explore. This is also borne out by additional analysis (Online Appendix Section 2), which shows that flat teams' relative performance improves when the individual exploration parameter (τ) is lower, whereas the crowd's advantage is most visible when the individual exploration parameter (τ) is higher.

We next try to understand how hierarchical teams produce similar and high-resolution beliefs by comparing the microstructure of flat and hierarchical teams. In Online Appendix Section 3, we examine the effect of symmetric and asymmetric influence within dyads. Dyadic asymmetry distinguishes the relationships between apex actors and other team members in hierarchical teams. More generally, asymmetric dyadic ties are the smallest building block of hierarchical structures (Figure 3.1 in the online appendix). Analysis of models with two agents shows that influence (both symmetric and asymmetric) increases similarity of beliefs across initially heterogeneous agents but also, lowers the resolution of beliefs (Figure 3.2 in the online appendix). The intuition can be seen by averaging two vectors of dissimilar beliefs with the same initial resolution. Although they become more similar after averaging, each vector must necessarily also have lower resolution if the elements with the strongest beliefs are not the same across agents prior to the influence process. Because more diffuse beliefs lead to more switching in actions, agents in flat and hierarchical teams that are subject to influence engage in more initial exploration than their counterparts in crowds.

Although asymmetry of influence does not directly produce greater similarity among agents' beliefs than a symmetric influence structure, it does make the resolution asymmetric between agents. Agents that influence others but are themselves resistant to influence maintain beliefs with higher resolution. This differential resolution triggers a well-known effect in the context of coupled learning: finding better combinations of actions is easier in coupled learning situations when uncoordinated simultaneous adjustment is avoided (Lounamaa and March 1987, Puranam and Swamy 2016). By staying relatively stable in their behavior, although broadcasting their beliefs, influential apex agents in hierarchical teams ensure that their favored options are visited by others and that the full potential of these options (and the benefits of coordination) is revealed. With high payoffs, this in turn leads to higher average resolution of all agents' beliefs, which further increases convergence and prevents excess exploration.

Dyadic asymmetry thus plays a key role in attaining both high similarity and high resolution and therefore, rapid convergence to actions with good (if not always the best) payoffs. This does not necessarily mean that any number of asymmetric ties will always create rapid convergence of actions, however. Note that in hierarchical teams as we have modeled them, all the ties of the apex agent are asymmetric, so that no agent influences the apex agent directly or through any other agent. This produces acyclicity for the apex agent. In Online Appendix Section 4, we show that teams with the same number of six asymmetric ties but randomly dispersed throughout the team do not create high-resolution beliefs that result in behavioral stability (Figure 4.1 in the online appendix). This is because hierarchical teams derive advantages at convergence not only from dyadic asymmetry but also, from acyclicity of influence around the leader. Disrupting acyclicity of influence in a hierarchical team by making the apex agent open to symmetric influence from even just one subordinate agent leads to a sharp drop off in performance (Figure 4.2 in the online appendix).

To summarize, we find that hierarchical influence structures differ from flat structures and unconnected crowds by achieving rapid performance improvements and convergence of behavior. This rapid convergence property is not by itself always beneficial, as it can lead to premature convergence to good but not optimal choices, a risk that is amplified when the typical nonglobal peak offers substantially lower payoffs than the global peak. However, rapid convergence is valuable in situations that require coordination where the variability of payoffs is modest and the environment is munificent (i.e., even the nonglobal peak payoffs are fairly attractive) and obviously, when speed matters. Hierarchical teams can dominate in such situations by producing highly similar beliefs that also attain greater resolution through its structural attributes of dyadic asymmetry and acyclicity. Agents in flat structures also develop highly similar beliefs, but because they do not feature differential resolution, they do not attain high average levels of resolution and therefore, enjoy the rapid convergence advantages of hierarchical teams. Agents in crowds, as they learn from the environment alone, develop high-resolution beliefs, but their beliefs may not be as similar (because they do not influence each other) and can rely only on the bottom-up convergence provided through feedback from the shared task environment. Given enough time and appropriate individual exploration, however, either flat teams or crowds can perform well without risking the fixation to inferior actions to which hierarchical teams are prone.

6. Exploring More Influence Structures

Our analysis shows that team structures, defined through a topology of ties and the weight that individual agents put on learning through environmental payoffs on their own actions versus learning socially from their connections (w), create particular distributions of beliefs (in terms of similarity and resolution) that in turn create particular tendencies for exploration and stability. Here and in Online Appendix Section 5, we explore the effect of additional team structures.

By changing w in a fully connected team, we achieve a continuous transition between crowds and flat teams (Figure 5.1 in the online appendix). This analysis shows us first that the optimal level of w in a fully connected team is quite high (0.9), indicating that it is valuable to maintain individual autonomy in these fully flat teams. Second, we see that individual autonomy is much more valuable when the cost of entrapment is higher (when the task environment has low δ_1). Third, we see that the benefit of individual autonomy over social influence is much more pronounced in the long run. This is because autonomous agents require more time to learn from the environment (recall that crowds outperformed other structures only in the long run).

We have shown that hierarchical teams introduce an element of stability to flat teams by adding an agent that is resistant to social influence. This agent learns from the environment (and therefore, indirectly from the actions of other agents) as much as the other agents (no more, no less) but differs from them in not accepting any social influence (in our baseline model of the hierarchical team). In Figure 5.2 in the online appendix, we examine the effect of varying the susceptibility of the apex agent to influence by varying its agent-specific w. This creates a smooth transition between models of flat and hierarchical teams. We find that in task environments where both coordination and search matter ($\chi = 1$), not only is the immunity to influence of the apex agent valuable (despite no knowledge advantage), the optimal apex agent is the one that is completely oblivious to social influence from other agents while still being receptive to feedback from the task environment (which does depend on the actions of other agents). In general, for hierarchical teams to do significantly worse when both search and coordination matter requires that we shut off learning from own experience for the apex agent—in other words, we have to impair the adaptive properties of the apex agent.

How does the susceptibility to social influence of the subordinates affect team performance? We address this question by varying the weights that socially influenced agents place on their own experience versus on the influence of other agents within different structures. We also introduce a star structure in this graph, where only the apex agent has ties to other agents, to consider a structure in which subordinates have no ties to each other, representing separation between subordinates that is quite likely in larger hierarchical structures. In Figure 5.3 in the online appendix, where $\chi = 1$, flat teams generally do worse than crowds and hierarchical teams. The star structure does best as it entails low mutual influence among the subordinates and therefore, less exploration. In Figure 5.4 in the online appendix, where $\chi = 0$, flat teams outperform hierarchical teams and crowds as long as agents retain significant autonomy ($w > 0.5$). When $w < 0.5$, flat teams do worse than other structures regardless of the level of χ, as high levels of social influence lead to downplaying what is learned from the environment.

7. Robustness Checks

In the models we have reported that interdependence, where it exists, is global—that is, everybody in the group needs to coordinate with everybody else. We could also specify interdependence between subsets of agents, such that there is a need for coordination only among agents connected by influence ties but not otherwise (e.g., Ethiraj and Levinthal 2004, Clement and Puranam 2018). In supplemental analyses shown in Online Appendix Section 6, we model a task environment where the task interdependencies between agents are not global but are isomorphic to the influence structure. Our results remain qualitatively unchanged.

We also consider how the effects of symmetric and hierarchical social influence may scale up. The hierarchical team has a single layer. Although teams can be adequately represented in this way, complex organizations likely require multiple layers. Similarly, the flat team is a particular form of symmetric structure with maximum density. It is unrealistic to expect larger organizations to have that same structure. In Online Appendix Section 7, we extend our analysis of hierarchical structures to a multilayered branching hierarchy and of flat structures to random graphs with symmetric dyadic ties (communities) that have the same density as that of the branching hierarchy. The crowd remains the baseline for comparison. We also allow for a form of reinforcement learning that accounts for the law of recency. Finally, we also explore much larger organizations than we have in the baseline analysis. We find that making the convergence problem easier by increasing the number of agents relative to the number of alternatives ($N >> S$) leads to an improvement in both crowd and community relative to hierarchical structures. In contrast, when S increases relative to N or both N and S scale up, hierarchy continues to do well. Thus, our conclusions about the advantage of hierarchical structures in task environments featuring both search and coordination remain qualitatively unchanged with scale.

8. Discussion and Conclusion

Organizational adaptation typically occurs in task environments that feature both variability in payoffs and interdependence. This requires finding a balance between sufficient exploration and rapid convergence to exploit the fruits of exploration. Crowds without mutual influence on beliefs only have access to the bottom-up convergence created by interdependence in the task environment. Social influence structures that affect beliefs—hierarchical as well as flat—facilitate both more top-down (imposed) exploration and more rapid convergence than crowds but in different proportions. Flat structures with symmetric influence stimulate greater exploration than hierarchies by lowering the resolution (sharpness) of beliefs of their members, whereas hierarchical structures, by creating differential resolution in beliefs among agents, aid more rapid discovery and convergence to good if not optimal organizational actions than either flat teams or crowds. In other words, hierarchical influence balances individual exploration and group-level convergence, whereas nonhierarchical (symmetric) influence increases exploration (which is useful for search) but also results in lower-resolution beliefs (which are detrimental to coordination). It is when both a degree of exploration to find good organizational actions as well as rapid convergence are important that hierarchical structures of influence on beliefs shine.

Crucially, this benefit of hierarchical influence does not require the apex agent to have control over subordinates' actions, differential knowledge, distinctive decision rights, or any other differentiating quality. Instead, the balance between exploration and exploitation at the organizational level arises through a second-order division of labor in search between influential and influenced agents, which is strikingly different from that traditionally assumed in a command and control hierarchy, where the leaders search and the followers execute or the followers search, leaders approve, and the followers then execute (e.g., Rivkin and Siggelkow 2003). Rather, when organizational adaptation requires all members to contribute to the search for valuable interdependent actions, leaders may provide stability, whereas followers produce the variation needed for search.

In agile development teams, for instance, our results underline the value of a hierarchical influence structure given that both search and coordination are typically important for such teams. To be precise, although the team still benefits from being shielded from the interference of hierarchical elements from outside the team (e.g., such as reporting, accountability, or need for formal approval or adherence to a centralized plan), it can in fact benefit from a source of asymmetric influence *within* the team. This is consistent with what Dönmez et al. (2016) find with respect to the critical role of project owners in creating stability during agile development. Our results suggest that although scrum masters, by maintaining the symmetry of relationships between developers, may deliver exploration, project owners, as apex actors, can create the convergence that is required to stop excessive search. Our results also suggest that recognizing and preserving these distinct functions of product owner and scrum master roles are important. Further and perhaps most importantly, the rapid iteration that makes agile development valuable (Ghosh and Wu 2018) is fully compatible with and even benefits from hierarchical influence, contrary to the rhetoric of flat management that has begun to be strongly associated with agile development as it has spread gradually beyond software development.

In thinking about the implications of our findings, it is worth emphasizing that the hierarchical structures we have considered are restricted to those that involve influence on beliefs. Formal authority hierarchies as found in a bureaucracy will typically bundle together influence over not only beliefs, but also actions through rewards and penalties as well as systems that differentially allocate decision rights and produce differential expertise at various levels. Existing theoretical analysis shows that hierarchies adapt better than other organizational forms when the agents who can identify superior solutions in a complex environment are matched to problems appropriately (Radner 1992, Garicano 2000; also see Rivkin and Siggelkow 2003, Siggelkow and Rivkin 2005) or can control the actions of others in a manner that mitigates conflicts of interest (Milgrom and Roberts 1988, Williamson 1991). Our analysis shows that hierarchies can be adaptive even *without* such effects on actions merely by creating variation and similarity in beliefs in a manner that neither crowds nor flat structures can achieve—a fact of particular relevance in knowledge-based production settings. These findings may help explain why hierarchical structures continue to flourish in innovation-intensive sectors (Lee and Edmondson 2017, Freeland and Zuckerman 2018) and why even those organizations that have dispensed with traditional administrative hierarchy retain hierarchical influence structures (Laloux 2014, Bunderson and Sanner 2017).

Our findings also imply that informal status differences, which produce asymmetric influence on beliefs, may have adaptive advantages for groups that face both search and coordination challenges. Perhaps this is a reason for status differences to be ubiquitous and for them to emerge spontaneously in groups (Rosa and Mazur 1979, Ridgeway 2019). A similar functional argument has been made by others who have shown that status can be a coordination device (e.g., Mark 2018). Experiments with dyads have shown that status can facilitate coordination by creating a salient focal point (de Kwaadsteniet and van Dijk 2010). In larger groups too, status differences facilitate coordination, as when higher-status actors contribute to collective action earlier and in a greater amount, prompting more contributions from lower-status actors (Kumru and Vesterlund 2010, Simpson et al. 2012). Our study shows that high-status individuals influencing low-status individuals' beliefs can aid group-level adaptation more effectively than egalitarian or unconnected groups even when adaptation requires search, not just coordination.

The distinctive features of organizational adaptation our model highlights also point to an important difference from social diffusion and therefore, to boundary conditions on the application of our results to social diffusion at large and conversely, on the social diffusion results to organizational contexts. Influence on one's own beliefs (and therefore, actions) can arise via the beliefs, actions, or outcomes of the other. However, this can unfold very differently outside versus inside organizations. Although actions or outcomes may be observed outside organizations, beliefs are typically not accessible. Individuals are constrained to observing each other's actions and inferring beliefs (and thus, future behavior) from those actions. In situations where agents lack the means to communicate about their beliefs or require social affirmation from observed actions, models in which agents only learn from each other's top choice (as revealed in their "vote" or action) are appropriate and may be sufficient (e.g., Centola and Macy 2007, Guilbeault et al. 2018).

In contrast, learning about other's beliefs requires communication and in some cases, the ability to trust what is communicated. Formal organizations are uniquely suited to satisfy these conditions, allowing individuals to communicate about decision premises—a key and distinguishing feature of influence processes within organizations (Simon 1947, 1981; Argote 2012; Baum et al. 2000; Reagans and McEvily 2003). Meanwhile, some relevant actions and outcomes may not be mutually observable in organizations because of division of labor and specialization, making belief exchange (which can occur even when actors operate in distinct action spaces) the only feasible path to social learning (e.g., see the literature on communication constraints across specialists cited in Knudsen and Srikanth 2014 and Puranam and Swamy 2016). Accordingly, many prior studies of organizational search have modeled belief exchange, including March (1991), Miller et al. (2006), and Knudsen and Srikanth (2014).

This study is unique in examining the effect of belief exchange on adaptation in task environments that combine the dual challenges of search and coordination. The joint importance of search and coordination, as we have argued, is pervasive in organizational adaptation. Even processes that can be modeled and created in the laboratory as pure coordination dynamics, such as the emergence of communication conventions (e.g., Steels 2001, Centola and Baronchelli 2015, Spike et al. 2017), can involve an element of search when created against a background of preexisting conventions (e.g., Fay et al. 2008, Guilbeault et al. 2021) and create frictions because of the need for unlearning (Koçak and Puranam 2022). Conversely, processes of search that do not require any coordination when carried outside organizations may require coordination when carried by agents in organizations whose choices are interdependent because of complementarities or common constraints such as budgets.

Influence on beliefs within organizations is also often heterogeneous (Mason et al. 2007) because of asymmetric and acyclic structures. More broadly, organizational structures may be seen as an effort to modify the direction and strength of influence that is ubiquitous in human groups (e.g., Salganik and Watts 2008, Muchnik et al. 2013). Formal hierarchy is a particular solution, with influence flowing, on balance, in only one direction. This allows hierarchies to present a single focal point for coordination—the leader (Foss 2001)—unlike flat structures like communities that prompt similarity but not necessarily a rapid convergence in action when actors are also interdependent. Many alternative forms can be created between the ideal types of hierarchies and flat structures, and indeed, teams in organizations are found to approximate these to various degrees (Bunderson et al. 2016, Wellman et al. 2019). Our results on the components of hierarchical influence show that dyadic asymmetry and acyclicity both contribute to give hierarchy its advantages at rapid convergence to good (if not optimal) organizational actions when the task environment features interdependencies.

We highlight that our results also show that hierarchical influence is not universally superior. Michel's "iron law" of oligarchy (from 1911) may still have some force in the knowledge economy, but its applicability certainly has recognizable limits. Even in contexts where search and coordination are both necessary for adaptation, crowds can adapt well in the long run as long there is some individual-level exploration. These findings respond to the contention of Piezunka et al. (2020) that work on the wisdom of crowds, which privileges the diversity of static prior beliefs, should be supplemented with a consideration of learning dynamics. In our models, problems of organizational adaptation (i.e., which entail both search and coordination) can be solved through bottom-up self-organizing processes, in which crowds of actors implicitly adapt to each other by learning from correlated feedback. Given some level of individual exploration, the crowd can outperform both the flat and hierarchical team even in environments that require both search and coordination. However, this requires time. This supports the notion that performance incentives may be a better tool to support coordinated search than social influence when selection acts slowly (for example, for the development of a new dominant design in contrast to the development of a specific new product).

In this work, we have not addressed issues of incentive conflict. We assumed that payoffs to picking compatible actions are positively correlated across agents because each agent's payoff is positively correlated with the organizational payoff (because we assume the agency problem has been sufficiently resolved to induce such a positive correlation). Future studies may consider how incentive misalignment interacts with differences in influence. Intuitively, hierarchical influence would play an even greater role under misalignment of interests, as it would encourage alignment of behavior toward the apex actor's beliefs even when the direct payoffs to the agents would suggest otherwise. Further, the stability of behaviour that accounts for the advantage of hierarchical influence may be disrupted if superiors are given incentives that make them more exploratory. Conversely, an accepted status hierarchy that makes leaders secure in their positions may encourage even more (perhaps too much) stability of actions.

Our model is a reminder of the importance of the early experimental work on the efficacy of teams with different communication structures in a context where no member has wisdom or special insight into the problem at hand (Bavelas 1950, Leavitt 1951). Even though the task used in those experiments (an "offline" search problem of aggregating dispersed information that is already present within the team) is fundamentally different from the task that we have modeled in this paper (an "online" search problem where members learn from the environment while solving a particular problem), the findings presented in those studies similarly point to the impact of communication structure being contingent on type of task; although centralized groups in those experiments were the most successful in aggregating distributed private information, the decentralized groups appeared to be better at generating and utilizing a novel method for solving the problem. The predictions that our models produce can be tested with further experimental studies. For instance, one might design laboratory experiments modifying those used by Mason and Watts (2012), Shore et al. (2015), or Acerbi et al. (2016) to incorporate influence on beliefs that is either symmetric or asymmetric. Further, one could alter the extent of the variability of the payoffs (as Mason and Watts 2012 do) and also, of interdependence between actors (which they do not).

In sum, this paper contributes to the broader research on comparative organizational forms (e.g., Williamson 1991, Freeland and Zuckerman 2018, Bremner and Eisenhardt 2021) by showing that hierarchical and nonhierarchical influence structures produce reliably different outcomes even when the difference between these forms is reduced to patterns of social influence on beliefs. Further, it contributes to research on functions of hierarchy (e.g., Anderson and Willer 2014, Bunderson et al. 2016, Tarakci et al. 2016, Greer et al. 2018) by showing that hierarchies of influence can affect organizational outcomes even without correlated expertise, authority, or power. It suggests new avenues of research on how to design hierarchical structures by examining the differential effects of its constituent elements (dyadic asymmetry and acyclicity). It also extends research on adaptive properties of organizational forms (e.g., March 1991, Siggelkow and Rivkin 2005, Lazer and Friedman 2007, Fang et al. 2010) by introducing coordination as an important element for adaptation in some task environments and developing a new modeling framework to separately tune the challenges of search and coordination and to study how a variety of influence structures performs in such environments.

Acknowledgments
Ekin İlseven assisted with visualization of results. Vikas Aggarwal, Stefano Brusoni, Julien Clement, Martin Gargiulo, Maciej Workiewicz, Andy Wu, Senior Editor Sameer Srivastava, and anonymous reviewers provided valuable suggestions on earlier versions of the manuscript. Participants of talks at the Carnegie School of Organizational Learning Conference (2018), Santa Fe Institute, Institut Européen d'Administration des Affaires, the University of Texas Austin, Duke, the University of North Carolina at Chapel Hill, Bocconi, and Northwestern gave helpful comments.

Endnotes

[1] The term hierarchy has variously been used to refer to uneven distribution of a valued resource (Bunderson et al. 2016, Greer et al. 2018), such as pay (Halevy et al. 2012), expertise (e.g., Garicano 2000, Tarakci et al. 2016), power, or status (Martin 2009, Gruenfeld and Tiedens 2010, Simpson et al. 2012). It may also refer to a nested containment structure (Simon 1962, Levinthal and Workiewicz 2018). Further, a hierarchy of authority itself may bundle multiple asymmetric relationships, such as directed influence over beliefs, a distribution of decision rights, rewards, and penalties (Coase 1937, Williamson 1991). Our focus is exclusively on a hierarchical structure of influence over beliefs.

[2] This is in contrast with agency theoretic approaches that focus on how to make the correlation between individual and organizational payoffs sufficiently positive given that agents allocate effort across tasks with (at least probabilistically) known returns to maximize their private benefits (e.g., Holmstrom 1982).

[3] We could specify a profile of compatible but different actions across agents; however, doing so would not alter the model mechanics as long as the payoffs across agents are aligned (i.e., the payoffs to picking compatible actions are positively correlated across agents, which will be true if each agent's payoff is positively correlated with the organizational payoff).

[4] In contexts where communication is constrained or may be deceptive, agents may have to infer or see proof of beliefs in the form of actions if that is feasible. In those cases, simple or complex contagion might be a more appropriate model of social influence (Centola and Macy 2007).

[5] Our use of the term crowd in reference to a collection of independent individuals who exchange no information in the decision-making process is more restrictive than the usage in economic sociology (e.g., trading crowds in Baker 1984) or in references to crowdfunding, where there is no assumption of lack of communication—only that there is no direct control of behavior. Our usage, however, is broader than that in the wisdom of crowds literature, where there is the additional assumption of the absence of task interdependence. The wisdom of crowds (Surowiecki 2005) refers to the improvement in accuracy of *static* beliefs by pooling them through averaging and making the group decision based on this averaged belief. In contrast, our analysis of crowds examines average performance of individuals engaged in *dynamic* learning processes. The performance of crowds, as sets of individuals who inhabit structures of no influence, is compared with those in structures with symmetric or hierarchic social influence as they learn in task environments that feature varying degrees of interdependence.

[6] To be precise, given agent-level exploration ($\tau > 0$), the phenomenon of trapping is better described as being "stuck long enough to lower cumulative performance at time T" rather than being trapped forever. Prior research has shown that interdependence among choices can create the risk of trapping on poor combinations of choices for a myopically searching actor (Levinthal 1997). In our setting, the choices are made by different agents, and the scope of search is the entire task environment. However, interdependence between actors creates an analogous risk of trapping on poor group-level choices not because of local search but because of the gains from coordination created by interdependence. This produces the intertemporal trade-off to rapid convergence that seems to give hierarchical teams an advantage that decays over time in task environments where both search and coordination matter.

References

Acerbi A, Tennie C, Mesoudi A (2016) Social learning solves the problem of narrow-peaked search landscapes: Experimental evidence in humans. *Roy. Soc. Open Sci.* 3(9):160215.

Aggarwal VA, Posen HE, Workiewicz M (2017) Adaptive capacity to technological change: A microfoundational approach. *Strategic Management J.* 38(6):1212–1231.

Ahl V, Allen TF (1996) Hierarchy theory: A vision, vocabulary, and epistemology (Columbia University Press, New York).

Anderson C, Willer R (2014) Do status hierarchies benefit groups? A bounded functionalist account of status. *The Psychology of Social Status* (Springer, New York), 47–70.

Argote L (2012) *Organizational Learning: Creating, Retaining and Transferring Knowledge* (Springer Science & Business Media, New York).

Arrow KJ (1974) *The Limits of Organization* (WW Norton & Company, New York).

Baker WE (1984) The social structure of a national securities market. *Amer. J. Sociol.* 89(4):775–811.

Bala V, Goyal S (1998) Learning from neighbours. *Rev. Econom. Stud.* 65(3):595–621.

Baum JA, Li SX, Usher JM (2000) Making the next move: How experiential and vicarious learning shape the locations of chains' acquisitions. *Admin. Sci. Quart.* 45(4):766–801.

Bavelas A (1950) Communication patterns in task-oriented groups. *J. Acoustical Soc. Amer.* 22(6):725–730.

Becker J, Brackbill D, Centola D (2017) Network dynamics of social influence in the wisdom of crowds. *Proc. Natl. Acad. Sci. USA* 114(26):E5070–E5076.

Blank S (2013) Why the lean start-up changes everything. *Harvard Bus. Rev.* 91(5):63–72.

Brackbill D, Centola D (2020) Impact of network structure on collective learning: An experimental study in a data science competition. *PLoS One* 15(9):e0237978.

Bremner RP, Eisenhardt KM (2021) Organizing form, experimentation, and performance: Innovation in the nascent civilian drone industry. *Organ. Sci.*, ePub ahead of print December 8, https://doi.org/10.1287/orsc.2021.1505.

Bunderson JS, Sanner B (2017). How and when can social hierarchy promote learning in groups? *The Oxford Handbook of Group and Organizational Learning.*

Bunderson JS, Van Der Vegt GS, Cantimur Y, Rink F (2016) Different views of hierarchy and why they matter: Hierarchy as inequality or as cascading influence. *Acad. Management J.* 59(4):1265–1289.

Bush RR, Mosteller F (1955) Stochastic models for learning (John Wiley & Sons, Inc., Oxford, England).

Centola D (2015) The social origins of networks and diffusion. *Amer. J. Sociol.* 120(5):1295–1338.

Centola D, Baronchelli A (2015) The spontaneous emergence of conventions: An experimental study of cultural evolution. *Proc. Natl. Acad. Sci. USA* 112(7):1989–1994.

Centola D, Macy M (2007) Complex contagions and the weakness of long ties. *Amer. J. Sociol.* 113(3):702–734.

Clement J, Puranam P (2018) Searching for structure: Formal organization design as a guide to network evolution. *Management Sci.* 64(8):3879–3895.

Chandler AD (1990) *Strategy and Structure: Chapters in the History of the Industrial Enterprise*, vol. 120 (MIT Press, Cambridge, MA).

Coase RH (1937) The nature of the firm. *Economica* 4(16):386–405.

De Groot MH (1974) Reaching a consensus. *J. Amer. Statist. Assoc.* 69(345):118–121.

de Kwaadsteniet EW, van Dijk E (2010) Social status as a cue for tacit coordination. *J. Experimental Soc. Psych.* 46(3):515–524.

DellaPosta D, Shi Y, Macy M (2015) Why do liberals drink lattes? *Amer. J. Sociol.* 120(5):1473–1511.

Denning S (2018) *The Age of Agile: How Smart Companies Are Transforming the Way Work Gets Done* (Amacom, New York).

Dönmez D, Grote G, Brusoni S (2016) Routine interdependencies as a source of stability and flexibility. A study of agile software development teams. *Inform. Organ.* 26(3):63–83.

Erev I, Roth AE (1998) Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *Amer. Econom. Rev.* 848–881.

Ethiraj SK, Levinthal D (2004) Modularity and innovation in complex systems. *Management Sci.* 50(2):159–173.

Fang C, Lee J, Schilling MA (2010) Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organ. Sci.* 21(3):625–642.

Fay N, Garrod S, Roberts L (2008) The fitness and functionality of culturally evolved communication systems. *Philos. Trans. Roy. Soc. London B Biological Sci.* 363(1509):3553–3561.

Foss NJ (2001) Leadership, beliefs and coordination: An explorative discussion. *Industrial Corporate Change* 10(2):357–388.

Freeland RF, Zuckerman EW (2018) The problems and promise of hierarchy: Voice rights and the firm. *Sociol. Sci.* 5(7):143–181.

Friedkin NE (2006) *A Structural Theory of Social Influence*, vol. 13 (Cambridge University Press, Cambridge, UK).

Friedkin NE, Johnsen EC (1990) Social influence and opinions. *J. Math. Sociol.* 15(3-4):193–206.

Friedkin NE, Johnsen EC (2002) Control loss and Fayol's gangplanks. *Soc. Networks* 24(4):395–406.

Friedkin NE, Proskurnikov AV, Mei W, Bullo F (2019) Mathematical structures in group decision-making on resource allocation distributions. *Sci. Rep.* 9(1):1377.

Galbraith JR (1973) *Designing Complex Organizations* (Addison-Wesley Longman Publishing Co., Inc., Boston).

Garicano L (2000) Hierarchies and the organization of knowledge in production. *J. Political Econom.* 108(5):874–904.

Ghosh S, Wu A (2018) Iterative coordination in search. *Acad. Management Proc.* 2018(1):13524.

Greer LL, de Jong BA, Schouten ME, Dannals JE (2018) Why and when hierarchy impacts team effectiveness: A meta-analytic integration. *J. Appl. Psych.* 103(6):591–613.

Gruenfeld DH, Tiedens LZ (2010) Organizational preferences and their consequences. *Handbook of Social Psychology*, 1252–1287.

Guilbeault D, Baronchelli A, Centola D (2021) Experimental evidence for scale-induced category convergence across populations. *Nature Comm.* 12(1):327.

Guilbeault D, Becker J, Centola D (2018). Complex contagions: A decade in review. *Complex Spreading Phenomena in Social Systems*, 3–25.

Guinan PJ, Parise S, Langowitz N (2019) Creating an innovative digital project team: Levers to enable digital transformation. *Bus. Horizons* 62(6):717–727.

Halevy N, Chou EY, Galinsky AD, Murnighan JK (2012) When hierarchy wins: Evidence from the national basketball association. *Soc. Psych. Personality Sci.* 3(4):398–406.

Henderson RM, Clark KB (1990) Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. *Admin. Sci. Quart.* 35(1):9–30.

Hill LA (2020) Being the agile boss. *MIT Sloan Management Rev.* 62(1):7–10.

Holland JH (1975) *Adaptation in Natural and Artificial Systems* (University of Michigan Press, Ann Arbor, MI).

Holmstrom B (1982) Moral hazard in teams. *Bell J. Econom.* 13(2):324–340.

Jensen MC, Meckling WH (1976) Theory of the firm: Managerial behavior, agency costs and ownership structure. *J. Financial Econom.* 3(4):305–360.

Kennedy J, Mendes R (2002) Population structure and particle swarm performance. *Proc. 2002 Congress Evolutionary Comput. 2002 CEC'02*, vol. 2 (IEEE), 1671–1676.

Keum DD, See KE (2017) The influence of hierarchy on idea generation and selection in the innovation process. *Organ. Sci.* 28(4):653–669.

Knudsen T, Srikanth K (2014) Coordinated exploration: Organizing joint search by multiple specialists to overcome mutual confusion and joint myopia. *Admin. Sci. Quart.* 59(3):409–441.

Koçak Ö, Puranam P (2022) Separated by a common language: How the nature of code differences shapes communication success and code convergence *Management Sci*. Forthcoming.

Krackhardt D (1994) Graph theoretical dimensions of informal organizations. Carley KM, Prietula MJ, eds. *Computational Organization Theory* (Lawrence Erlbaum Associates, Inc., Hillsdale, NJ), 89–111.

Kumru CS, Vesterlund L (2010) The effect of status on charitable giving. *J. Public Econom. Theory* 12(4):709–735.

Laloux F (2014) *Reinventing Organizations: A Guide to Creating Organizations Inspired by the Next Stage in Human Consciousness* (Nelson Parker, Brussels).

Lave CA, March JG (1975) *An Introduction to Models in the Social Sciences* (University Press of America, Lanham, MD).

Lazer D, Friedman A (2007) The network structure of exploration and exploitation. *Admin. Sci. Quart.* 52(4):667–694.

Leavitt HJ (1951) Some effects of certain communication patterns on group performance. *J. Abnormal Psych.* 46(1):38–50.

Lee MY, Edmondson AC (2017) Self-managing organizations: Exploring the limits of less-hierarchical organizing. *Res. Organ. Behav.* 37:35–58.

Levinthal DA (1997) Adaptation on rugged landscapes. *Management Sci.* 43(7):934–950.

Levinthal DA, Posen HE (2007) Myopia of selection: Does organizational adaptation limit the efficacy of population selection? *Admin. Sci. Quart.* 52(4):586–620.

Levinthal DA, Workiewicz M (2018) When two bosses are better than one: Nearly decomposable systems and organizational adaptation. *Organ. Sci.* 29(2):207–224.

List C, Elsholtz C, Seeley TD (2009) Independence and interdependence in collective decision making: An agent-based model of nest-site choice by honeybee swarms. *Philos. Trans. Roy. Soc. London B Biological Sci.* 364(1518):755–762.

Lounamaa PH, March JG (1987) Adaptive coordination of a learning team. *Management Sci.* 33(1):107–123.

Luce RD (1959) On the possible psychophysical laws. *Psych. Rev.* 66(2):81–95.

March JG (1991) Exploration and exploitation in organizational learning. *Organ. Sci.* 2(1):71–87.

March JG, Levinthal DA (1993) The myopia of learning. *Strategic Management J.* 14(S2):95–112.

March JG, Simon HA (1958) *Organizations* (Wiley, New York).

Mark NP (2018) Status organizes cooperation: An evolutionary theory of status and social order. *Amer. J. Sociol.* 123(6):1601–1634.

Martin JL (2009) *Social Structures* (Princeton University Press, Princeton, NJ).

Mason W, Watts DJ (2012) Collaborative learning in networks. *Proc. Natl. Acad. Sci. USA* 109(3):764–769.

Mason WA, Conrey FR, Smith ER (2007) Situating social influence processes: Dynamic, multidirectional flows of influence within social networks. *Personality Soc. Psych. Rev.* 11(3):279–300.

Mihm J, Loch CH, Wilkinson D, Huberman BA (2010) Hierarchical structure and search in complex organizations. *Management Sci.* 56(5):831–848.

Milgrom P, Roberts J (1988) An economic approach to influence activities in organizations. *Amer. J. Sociol.* 94:S154–S179.

Miller JH, Page SE (2009) *Complex Adaptive Systems: An Introduction to Computational Models of Social Life* (Princeton University Press, Princeton, NJ).

Miller KD, Zhao M, Calantone RJ (2006) Adding interpersonal learning and tacit knowledge to March's exploration-exploitation model. *Acad. Management J.* 49(4):709–722.

Mintzberg H (1979) *The Structuring of Organizations* (Prentice Hall, Englewood Cliffs, NJ).

Muchnik L, Aral S, Taylor SJ (2013) Social influence bias: A randomized experiment. *Science* 341(6146):647–651.

Nelson R, Winter S (1982) *An Evolutionary Theory of the Firm* (Harvard University Press, Cambridge, MA).

Page S (2010) *Diversity and Complexity* (Princeton University Press, Princeton, NJ).

Piezunka H, Aggarwal VA, Posen HE (2020) Learning-by-participating: The dual role of structure in aggregating information and shaping learning. Preprint, submitted May 16, http://dx.doi.org/10.2139/ssrn.3425696.

Posen HE, Levinthal DA (2012) Chasing a moving target: Exploitation and exploration in dynamic environments. *Management Sci.* 58(3):587–601.

Puranam P (2018) *The Microstructure of Organizations* (Oxford University Press, Oxford, UK).

Puranam P, Håkonsson DD (2015) Valve's way. *J. Organ. Design* 4(2):2–4.

Puranam P, Swamy M (2016) How initial representations shape coupled learning processes. *Organ. Sci.* 27(2):323–335.

Puranam P, Raveendran M, Knudsen T (2012) Organization design: The epistemic interdependence perspective. *Acad. Management Rev.* 37(3):419–440.

Puranam P, Stieglitz N, Osman M, Pillutla MM (2015) Modelling bounded rationality in organizations: Progress and prospects. *Acad. Management Ann.* 9(1):337–392.

Radner R (1992) Hierarchy: The economics of managing. *J. Econom. Literature* 30(3):1382–1415.

Reagans R, McEvily B (2003) Network structure and knowledge transfer: The effects of cohesion and range. *Admin. Sci. Quart.* 48(2):240–267.

Ridgeway CL (2019) Status: Why is it everywhere? Why does it matter? (Russell Sage Foundation).

Rigby DK, Sutherland J, Takeuchi H (2016) Embracing agile. *Harvard Bus. Rev.* 94(5):40–50.

Rivkin JW, Siggelkow N (2003) Balancing search and stability: Interdependencies among elements of organizational design. *Management Sci.* 49(3):290–311.

Rosa E, Mazur A (1979) Incipient status in small groups. *Soc. Forces* 58(1):18–37.

Salganik MJ, Watts DJ (2008) Leading the herd astray: An experimental study of self-fulfilling prophecies in an artificial cultural market. *Soc. Psych. Quart.* 71(4):338–355.

Schelling TC (1960) *The Strategy of Conflict* (Harvard University Press, Cambridge, MA).

Schilling MA, Fang C (2014) When hubs forget, lie, and play favorites: Interpersonal network structure, information distortion, and organizational learning. *Strategic Management J.* 35(7):974–994.

Schwaber K, Sutherland J (2020) The Scrum Guide (Scrum Alliance). Accessed May 25, 2021, https://scrumguides.org/scrum-guide.html.

Shore J, Bernstein E, Lazer D (2015) Facts and figuring: An experimental investigation of network structure and performance in information and solution spaces. *Organ. Sci.* 26(5):1432–1446.

Siggelkow N, Rivkin JW (2005) Speed and search: Designing organizations for turbulence and complexity. *Organ. Sci.* 16(2):101–122.

Simon HA (1947) *Administrative Behavior, a Story of Decision Processes in Business Organization* (Macmillan, London).

Simon HA (1962) Architecture of complexity. *Proc. Amer. Philos. Soc.* 106(6):467–482.

Simon HA (1981) *The Sciences of the Artificial* (MIT Press, Cambridge, MA).

Simpson B, Willer R, Ridgeway CL (2012) Status hierarchies and the organization of collective action. *Sociol. Theory* 30(3):149–166.

Spike M, Stadler K, Kirby S, Smith K (2017) Minimal requirements for the emergence of learned signaling. *Cognitive Sci.* 41(3): 623–658.

Steels L (2001) Language games for autonomous robots. *IEEE Intelligent Systems* 16(5):16–22.

Surowiecki J (2005) *The Wisdom of Crowds* (Anchor, New York).

Sutton RS, Barto AG (1998) *Introduction to Reinforcement Learning*, vol. 135 (MIT Press, Cambridge, MA).

Tarakci M, Greer LL, Groenen PJF (2016) When does power disparity help or hurt group performance? *J. Appl. Psych.* 101(3): 415–429.

Thompson JD (1967) *Organizations in Action: Social Science Bases of Administrative Theory* (McGraw Hill, New York).

Weick KE (1974) *The Social Psychology of Organizing* (Addison-Wesley Publishing Company, Boston).

Wellman N, Applegate JM, Harlow J, Johnston EW (2019) Beyond the pyramid: Alternative formal hierarchical structures and team performance. *Acad. Management J.* 63(4):997–1027.

Williamson OE (1975) *Markets and Hierarchies* (Free Press, New York).

Williamson OE (1991) Comparative economic organization: The analysis of discrete structural alternatives. *Admin. Sci. Quart.* 36(2): 269–296.

Wu S, Ghosh A (2021) Iterative coordination and innovation: Prioritizing value over novelty. *Organ. Sci.*, ePub ahead of print October 11, https://doi.org/10.1287/orsc.2021.1499.

Özgecan Koçak is associate professor of organization and management at Emory University's Goizueta Business School. She received her PhD from the Graduate School of Business, Stanford University. Her research focuses on how shared understandings (such as communication codes, categorization systems, and identity schema) emerge and shape behavior in organizations and markets.

Daniel A. Levinthal is the Reginald H. Jones Professor of Corporate Strategy at the Wharton School, University of Pennsylvania. Levinthal works on issues of organizational adaptation and industry evolution, particularly in the context of technological change.

Phanish Puranam is a professor of strategy at INSEAD. He received his PhD from the Wharton School, University of Pennsylvania. His research focuses on how organizations work and how we can make them work better. His current interests include nonhierarchical organizations, culture, and how intelligent algorithms shape organizations.