ECONSTOR
Make Your Publications Visible.

A Service of
ZBW
Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Anjos, Fernando; Reagans, Ray

Article

Networks in the balance: An agent-based model of optimal exploitation

Journal of Organization Design

Provided in Cooperation with:
Organizational Design Community (ODC), Aarhus

Suggested Citation: Anjos, Fernando; Reagans, Ray (2020) : Networks in the balance: An agent-based model of optimal exploitation, Journal of Organization Design, ISSN 2245-408X, Springer, Cham, Vol. 9, Iss. 1, pp. 1-26, https://doi.org/10.1186/s41469-020-00084-x

This Version is available at:
https://hdl.handle.net/10419/252170

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

https://creativecommons.org/licenses/by/4.0/

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

Anjos and Reagans *Journal of Organization Design* (2020) 9:20
https://doi.org/10.1186/s41469-020-00084-x

Journal of Organization Design
a SpringerOpen Journal

RESEARCH **Open Access**

Networks in the balance: an agent-based model of optimal exploitation

Fernando Anjos[1*] and Ray Reagans[2]

* Correspondence: fernando.anjos@novasbe.pt
[1]NOVA School of Business and Economics, Campus de Carcavelos, Rua da Holanda, n.1, 2775-405, Carcavelos, Portugal
Full list of author information is available at the end of the article

Abstract

Using an agent-based simulation, we illustrate how goal-seeking behavior affects network formation, learning, and performance. Our organization has one manager, who decides where to invest financial capital; individual workers, who decide where to work and prefer projects with larger budgets; and projects, which vary in quality. Our manager discovers high-quality projects from interactions with workers and allocates more capital to high-quality projects. When given an opportunity, our workers move to bigger-budget projects. We let our manager vary in terms of how much she exploits what she learns and allow our workers vary in terms of how sensitive they are to differences in capital. Our results highlight a contingency which shapes how goal-seeking behavior affects learning. The contingency is network fragility. Fragile connections decay quickly when individuals are not working together, while robust relationships decay more slowly. When relationships are robust, exploitation by our manager leads to a dense organizational network, improving information quality, and performance. Decisions by self-interested individuals (our manager and our workers) produce a virtuous learning cycle. When relationships are fragile, exploitation by our manager produces a sparse network, reducing information quality, and undermining performance. When network connections are fragile, the manager must find the right balance of exploitation and exploration, a balance which limits the rate at which workers move from one project to the next, allowing the manager to exploit some of what she knows, without undermining the very network which allows for useful information to be obtained.

Keywords: Exploitation, Exploration, Learning, Mobility, Organizational networks

Introduction

A basic idea in organizational theory is individuals can learn faster in an organization than alone (Argote and Ingram 2000). When individuals problem solve together, they can share what they are learning and potentially learn more quickly (Epple et al. 1996). The process is called mutual learning and while mutual learning can improve performance, it can also result in individuals focusing their problem solving efforts on improving sub-optimal ideas and solutions (Knudsen and Srikanth 2014). The value of mutual learning depends critically on how individuals problem solve and learn together (March 1991; Levinthal and March 1993). When an individual is working alone, he or she must strike a balance between improving existing ideas and solutions against

Springer Open

© The Author(s). 2020 **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

continuing to search for and experiment with alternative solutions to the assigned problem (Simon 1962, pp. 472–473; Sutton and Barto 1998; Denrell et al. 2004). The problem-solving modes are understood as exploitation and exploration. Each approach to problem solving is essential to learning, but there exists the potential for a tradeoff between the two. Time spent in one problem-solving mode comes at the expense of time spent in the alternative mode. In a seminal article, March (1991) describes a mutual learning process where an organization can realize the benefits of exploitation and exploration. Learning in an organization allows for a division of effort where individuals can employ different problem-solving modes to the benefit of the larger organization.

March's framework is a cornerstone in organizational learning research. An important element is often overlooked. The organization contains a network. March's organization contains a manager (the organizational code in March's language) who is connected to individual problem solvers. The connections allow the manager and the workers to adopt different approaches to learning and to share what they learn. Mutual learning is enhanced when the manager focuses on exploitation, while the workers focus on exploration. The relationships between the manager and individual workers allow March's manager to "exploit the explorations" of the individual workers (Fu 2020).

An organizational network provides conduits for knowledge transfer. The structure of the network affects the approach to learning that is more likely to dominate the mutual learning process (Lazer and Friedman 2007; Fang et al. 2010; Csaszar and Siggelkow 2010; Mason and Watts 2012; Schilling and Fang 2014; Boudreau and Lakhani 2015; Shore et al. 2015). When a knowledge source shares an idea with a recipient, the network structure affects how quickly the idea spreads from the original source to multiple recipients. When ideas diffuse rapidly, exploitation will be the dominant problem solving mode. If ideas diffuse slowly, the individuals are more likely to consider a wider array of solutions. A slow diffusion rate maintains system-level exploration. Indeed, the structure of the network could encourage groups of individuals to focus on improving different solutions, allowing the organization to realize the benefits of exploitation and exploration. Thus, existing research on mutual learning not only illustrates the need to balance exploitation and exploration; it also highlights how the broader organizational network affects the balance between the two.

Scholars have begun to consider mutual learning in dynamic networks (Clement and Puranam 2018; Songhori and García-Díaz 2018). Individuals enter and exit network connections out of a desire to improve performance. While we see considerable merit in this emerging line of research, we note a potential dilemma. When individuals exploit what they know to develop relationships, their choices could have the unintended consequence of reintroducing a tradeoff between exploitation and exploration. Indeed, the individuals could build a network that allows them to exploit what they know. But the network they build could also make it harder to identify better ideas and solutions. By exploiting what they know, the individuals could undermine the very network which allowed them to identify valuable ideas and solutions in the first place. Instead of an immediate tradeoff between exploitation and exploration, the tradeoff would occur between exploitation and the organizational network which facilitates exploration in the future.

We use an agent-based simulation to illustrate the indirect tradeoff between exploitation and exploration. Our organization contains a manager, individual workers, and projects. The three elements allow us to create a mutual learning process between the manager and the individual workers, with the individual choices influencing the decisions made by other actors, network formation, and ultimately organizational performance. We describe the framework in detail later in the manuscript. Sufficient to say for now, the framework allows us to analyze how exploitation affects network formation and how the resulting network affects firm performance. The simulation results illustrate the indirect tradeoff between exploitation and exploration but also illustrate conditions which allow for the tradeoff to be lessened and avoided altogether. The rest of the manuscript is as follows. We start with a discussion of mutual learning in organizations, beginning with March and then shifting to more recent research where the organizational network affects the problem-solving mode individuals are more likely to employ. Next, we describe why allowing for an endogenous network can reintroduce a tradeoff between exploitation and exploration, albeit indirectly through the effect exploitation can have on the broader organizational network.

Literature review

Prior research has established the positive effect knowledge transfer can have on organizational learning (Epple et al. 1991; Ingram and Simons 2002; Zuckerman and Sgourev 2006; Diwas et al. 2013). The positive effect is assumed to be a byproduct of a mutual learning process. The importance of mutual learning for organizational learning was first articulated by March (1991). March's organization contains a manager (March's organizational code) and individual problem solvers. The organization is given a problem to solve and there exists an array of potential solutions. Each solution is characterized by a fixed number of elements, with a choice for each element. The choices define choice sets and initially the performance implications of the different choice sets are unknown. It is unclear to the manager and the individual problem solvers what solution they should pursue. The members of the organization learn by doing. There is, however, an important division of labor in March's organization. The individual problem solvers work on different solutions and discover their performance implications. The manager learns from their current collective efforts and uses what he or she learns to direct their future problem-solving efforts. How much the manager learns from the individual problem solvers and how much the individuals learn from the manager defines the extent of mutual learning (March 1991, pp. 76–78). Organizational learning is enhanced when the manager is a "fast" learner, and the individual problem solvers are "slow" learners.

In more recent research, the individual problem solvers are connected and can share what they are learning directly with each other (Lazer and Friedman 2007; Fang et al. 2010; Csaszar and Siggelkow 2010; Mason and Watts 2012; Schilling and Fang 2014; Boudreau and Lakhani 2015; Shore et al. 2015). If an individual is connected to someone who has found a better solution to the assigned problem, the focal individual can adopt it. The broader network is important because the structure of the network affects the rate shared solutions diffuse across the organization. The rate affects the problem-solving mode that is more likely to dominate the mutual learning process. When an individual adopts a solution from a colleague, the amount of diversity in the solutions

being considered declines (Lazer and Friedman 2007, pp. 678–679; Fang et al. 2010, p. 632). The decline in diversity represents a drop in system-level exploration. The effect the decline in exploration has on performance depends on the kind of problems individuals are attempting to solve. Prior research has considered two kinds of problems, one where the choices in a choice set have independent effects on performance, and the other where the choices have interdependent effects on performance (Levinthal 1997, pp. 936–937). When choices have independent effects on performance, the problem is easier to solve. For each element of a solution, there is a single best choice. To find the best solution, the individual has to find the best choice for each element. The problem is more basic. When the choices have interdependent effects on performance, the problem is more challenging and it is harder to find the best solution. When the choices have interdependent performance implications, the value of a specific choice depends critically on the other choices an individual has made. Finding the right solutions requires considerably more search and experimentation. The problem is more complex.[1]

When individuals are solving a basic problem, the decline in system-level exploration that results from adopting superior solutions from colleagues is beneficial. When an individual adopts a superior solution from a colleague, the overall level of performance in the organization increases. The spread of a current best practice has a positive effect on current performance. The decision also has an effect on longer term performance. Since the choices have independent effects on performance, any solution which is superior to the recipient's current solution is also closer to the best solution. For a basic problem, the decline in system-level exploration that results from the spread of current best practices represents a culling of inferior solutions. The solutions that remain are closer to the best solution, where closeness is defined in terms of the number of choices that would need to change to turn a focal solution into the ideal solution. The spread of the current best practices not only improves the current level of performance but also provides a foundation for continued improvements.

However, when individuals are solving a complex problem, adopting a superior solution from a colleague will have a positive effect on current performance. But the decline in system-level exploration is likely to have a negative effect on longer-term performance. Since the choices have interdependent effects on performance, a solution with a lower level of performance could in fact be closer to the best solution. As a result, the spread of a current best practice is likely to also reduce experimentation with solutions that could have yielded superior performance in the future. For a complex problem, the spread of the current best practice will have a positive effect on immediate performance but a negative effect on longer-term learning outcomes, as individuals focus their problem-solving efforts on improving sub-optimal solutions.

[1]We use basic and complex to refer to problems with either a smooth or a rugged performance landscape (Kauffman 1993 Weinberger 1990; Levinthal 1997, pp. 936–937; Puranam et al. 2012, pp. 421–422). When the elements of the potential solutions to a problem have independent effects on performance, the performance landscape defined by the potential solutions is smooth. As interdependency increases, the landscape becomes more rugged. The critical issue is not simply interdependence but the stochastic nature of interdependence (Kauffman 1993, pp. 40–45; Rahmandad 2019). When a landscape is smooth, the problem is easier to solve. For example, the problem could be solved with selective trial and error experimentation. When the landscape is smooth, we say the problem is basic. When a landscape is rugged, considerably more search and experimentation is required to find the best solution. The problem is harder to solve and so we say the problem is complex.

The rate at which system-level exploration declines as individuals adopt solutions from their colleagues is a function of the broader organizational network. Diversity declines at a much faster rate in organizational networks that allow shared solutions to spread rapidly from one recipient to the next (Lazer and Friedman 2007, pp. 678–683; Fang et al. 2010, p. 632).[2] Thus, when individuals are solving a basic problem, organizational networks that allow for the rapid diffusion of knowledge are more beneficial. A "random" organizational network is an example of a network that facilitates rapid diffusion. The centralized network in March is another example. Basic problems benefit from more exploitation. However, when individuals are solving a more complex problem, organizational networks that slow the diffusion of knowledge are more beneficial. An example is a small world network (Watts and Strogatz 1998). A small world network contains a relatively large number of relationships inside each network community but relatively few connections between network communities. The structure of a small world network allows for rapid transmission within local network communities but limits the transfer of knowledge between communities. In a small world network, separate groups can focus on improving different solutions, allowing the organization to realize the benefits of exploitation and exploration.

In the research we have discussed thus far, the organizational network has either been fixed or mechanically determined. The individuals have little discretion in creating the organizational network. An emerging line of research considers mutual learning processes in more dynamic and endogenous networks (Schilling and Fang 2013; Clement and Puranam 2018; Songhori and García-Díaz 2018). Individuals are allowed to enter and exit network connections out of a desire to improve performance. Allowing network connections to be endogenous adds more "realism" to agent-based simulations. But when network connections are endogenous, a potential tradeoff between exploitation and exploration can re-emerge. When a network is endogenous and individuals have choice in developing relationships, one can imagine individuals will attempt to develop relationships they expect to improve their performance. If individuals can develop network connections with colleagues who currently have the "best" solutions and ideas, the network can become too myopic and undermine the network foundation for discovering new ideas and solutions. For example, if individuals working in a small-world network are allowed to develop relationships with superior performers, they could end up in a more random or even a more centralized network (Songhori and García-Díaz 2018). The change in the organizational network would improve current performance, but if the individuals are solving a complex problem, the increase in performance would come at the expense of the very network which allowed for the superior solutions to be identified.

The potential tradeoff has not been fully appreciated by scholars who study mutual learning in dynamic networks because while networks are allowed to change, the direction of change is often orthogonal to performance. For example, problem solvers could develop new network connections randomly (Schilling and Fang 2013). When the

[2]A related line of research also highlights the downside of rapid transmission, when individuals are not learning but simply exchanging knowledge and information (Cowan and Jonard 2004; Reagans and Zuckerman 2008). Networks that allow rapid transmission make it more likely that some individuals will accumulate whatever resources they need and stop trading, and when they stop trading, fewer people can accumulate all of the resources they need.

network formation process is influenced by performance expectations, the network formation process and performance are "loosely linked," which allows for performance-oriented changes to the organizational network to occur more slowly, increasing the odds the changes will ultimately be beneficial (Clement and Puranam 2018, pp. 3885–3887).

While not their research focus, Anjos and Reagans (2013) illustrate the potential tradeoff between exploitation and exploration. Anjos and Reagans use a simulation to analyze how different commitment strategies affect individual learning and performance. Individual agents possess specific capabilities. Some of the capabilities are more compatible and produce higher levels of performance. The individuals enter and exit collaborations as they attempt to identify the best combinations of capabilities. When agents attempt to fully exploit what they know, they increase the odds of coordination failures, which reduces the number of realized partnerships and ultimately their ability to learn.

Allowing individual choices to affect network formation adds realism to agent-based simulations. But the narrow focus on individual choices is also unrealistic. Individuals' choices and decisions are not made in a vacuum. Organizational features can influence what individuals decide to do. Considering the influence of individual choices in the context of an organization adds another layer of realism. For example, in Clement and Puranam, network choices are influenced by an individual's performance expectations and the organization's formal structure. In their framework, the formal structure influences individual choices and decisions, which in turn affects the network connections that develop. Our model is in this same spirit, although we focus on the importance of a different organizational feature. We model an organization which contains different projects. Our manager decides how much to invest in each project and our workers decide where to work. Those choices have important implications for network formation. Individuals who decide to work on the same projects are more likely to develop and maintain a relationship, which in turn affects what individuals learn and how well the firm performs.

The verbal model

We illustrate organizational learning with an agent-based simulation. The formal details of the model and simulation are discussed in the appendix. Currently, we describe the critical features of our framework. The organization contains one manager, a number of individual workers, and projects that vary in quality. The individuals work on different projects and economic value is created when the manager invests capital in high-quality projects and the individuals decide to work on those projects. Initially, the quality of the different projects is unknown. However, each project has an expert. The experts are informational seeds. The non-experts, which we refer to as individual workers, can learn a project's economic quality by working on the project with the expert or by being connected to someone who has worked with the expert. Our organization is characterized by mutual learning. The manager learns the quality of the different projects through his or her interactions with individual workers and the manager uses what he or she learns to decide how to allocate financial capital. Workers decide where to work and when choosing between two projects, they are more likely to select the project receiving more capital. Therefore, the manager can direct the efforts

of the individual workers through his or her financial allocation decisions. We allow the workers in our organization to vary in terms of how sensitive they are to differences in capital. The individual workers are mobile but they vary in terms of how sensitive they are to differences in budgets. When workers are less sensitive, they are less likely to be mobile. Less sensitive workers are slow-movers. When workers are more sensitive to differences in capital, they are more likely to be mobile and move between projects. More sensitive workers are fast-movers.[3]

The manager and the individual workers in our organization develop beliefs about the quality of the different projects and the beliefs they hold affect the choices they make. The beliefs correspond to subjective probabilities about project quality, with a probability for the project being high-quality and one minus that probability for project quality being low-quality. Our agents hold subjective beliefs about all of the projects. The beliefs represent stocks of knowledge which decay over time, unless replenished by experience. While the beliefs are subjective, more accurate beliefs result in better choices, which improve organizational performance. Accurate beliefs correspond to the economic reality of the firm. We define information quality as the difference between what individuals believe and the actual quality of the projects. Information quality improves with experience but it also depends critically on the broader organizational network. The economic quality of each project in our organization is subject to change (Anjos and Reagans 2013, pp. 7–8). Thus, the individuals must be able to update their beliefs as the environment changes. The network connections among members of the organization become critical.

The network connections in our organization are a byproduct of goal-oriented behavior. However, the individuals do not search for more productive interactions (Anjos and Reagans 2013; Clement and Puranam 2018; Songhori and García-Díaz 2018). Individuals want to work on "better" projects and network connections emerge as a byproduct of where individuals decide to work. We model network formation and decay as a function of the amount of time individuals spend working together on a project. Time spent working together is an organizational equivalent of physical proximity. Proximate individuals have more opportunities to develop and maintain a network connection (Small and Adler 2019). For example, Reagans and McEvily (2003, pp. 252–253) describe the positive effect working on the same projects can have on network formation (Hasan and Koning 2020). In our organization, a network connection develops as two individuals spend more time working together. When two individuals stop working together, their relationship is at risk for decay. Network connections are channels for information diffusion. In addition to their direct work experiences, the individual workers can learn project quality from each other. The individuals share the beliefs they hold about project quality, but it is important to emphasize that project quality is characterized by uncertainty. A project's quality transitions from high to low, and from low to high with some probability. The transitions create an environment characterized by uncertainty.

[3]Individual worker's knowledge (or knowledge stock) about a particular project starts to decay once they are disconnected from the project expert. What they know about a project decays and becomes more obsolete and imprecise when they are disconnected from the project expert. The decay rate is a parameter in our model.

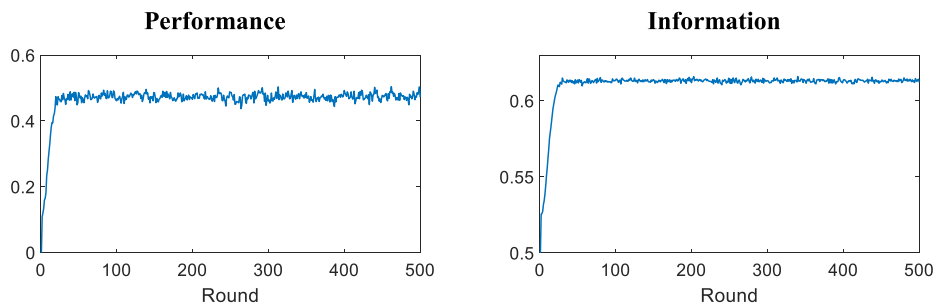

Fig. 1 The figure plots the evolution of organizational performance (left panel) and managerial information (right panel) across simulation rounds

It is important to be clear about how the diffusion process works in our framework. Recall, the experts are our informational seeds. An individual must either work with an expert or be connected to someone who has worked with the expert to have some insight into a project's quality. Our manager and individual workers update their beliefs as they interact and communicate with each other. When an individual is exposed to a range of beliefs about a specific project, he or she adopts the most extreme view, which could be his or her own view. The most extreme belief is closer to the truth. If the individuals in our organization were sharing quality signals which contained white noise, an individual would be better off taking a weighted average of the beliefs he or she hears. The weighted average would help cancel out white noise in the signal (Golub and Jackson 2010).

How the manager invests capital is the indicator of exploitation and exploration in our model. The manager operates between two allocation styles. At one extreme the manager "exploits" what he or she currently knows by investing in direct proportion to his or her expected level of project performance. With full exploitation, there is a one-to-one correspondence between expected quality and the amount of invested capital. At the other extreme, the manager "explores" the project space and invests capital in projects in an egalitarian fashion. With full exploration, there is no association between expected project performance the manager expects and how he or she allocates capital. Allocation decisions lying between the two extremes represent different combinations of exploitation and exploration.

Finally, it is important to emphasize that capital allocation is the only information channel from the manager to the individual workers. The interests between our manager and workers are compatible. Our manager wants to invest in high-quality projects and our workers want to work on projects with larger budgets. Thus, our workers have an incentive to help the manager identify high-quality projects and the manager has an incentive to invest more capital in those projects.

Simulation results

Figure 1 shows how our simulation unfolds over a path of 1000 rounds.[4] The initial network is empty and the economic quality of each project is assigned randomly. Individual workers are randomly assigned to projects. The left panel in Fig. 1 shows how

[4]To filter out noise, we repeat the simulation 1000 times and report time-t averages.

performance improves relative to the starting point. Performance improves quickly and stabilizes after a small number of rounds. The improvement in performance is due to the fact that over time, ties form, allowing the manager to have relevant information when making capital-allocation decisions. Improvement in information is visible in the panel on the right. Information quality refers to the accuracy of the manager's beliefs regarding project quality, measured as the distance between the manager's beliefs and the truth. Once the effect of initial conditions dissipates, the organization operates in an environment with relatively good information. We want to emphasize that our long-run outcomes refer to an "uncertain" steady state, where the manager needs to keep learning the quality of the projects at every point in time because project quality is subject to change. In other words, the environment can always change, and so some degree of learning is always required. Our analysis of organizational dynamics starts after round 500.

The results from our simulation are illustrated in Fig. 2. The horizontal axis is our exploitation parameter. The parameter describes how the manager allocates capital, given his or her beliefs about the economic quality of the different projects. The parameter varies from 0 to 1. The extreme of 0 corresponds to full exploration, where an equal amount of capital is allocated across projects; the extreme of 1 corresponds to full exploitation, where only the currently most promising project, from the manager's perspective, receives funding. The manager could decide to operate at one of the extremes, and the manager could also choose a capital allocation which is a combination of exploitation and exploration. Allocation amounts below .5 (50%) represent more exploration than exploitation, while amounts above 50% represent more exploitation than exploration. The vertical axis in the left panel is firm performance. Two lines are illustrated in the left panel of Fig. 2. For the dashed line, the individual workers are fast-movers. For the solid line, the individual workers are slow-movers. Recall that worker speed refers to how sensitive workers are to budgetary motives. A fast worker is easily lured into working on better-funded projects, even if the funding difference is small. A

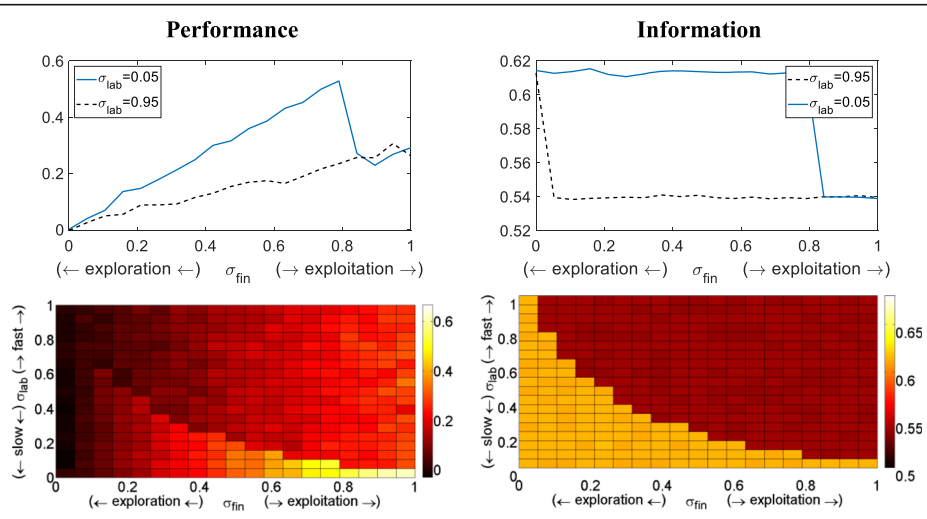

Fig. 2 The figure plots how performance (left) and information (right) vary as a function of the exploration-exploitation balance and labor speed. In the top panels labor is either slow (blue lines) or fast (dashed black lines). Bottom panels plot outcomes for the full parameter space

slow worker tends to stay with his or her default assignment (invariant over time), and only looks to move to another project if capital allocations are significantly different. It is useful to start by contrasting the performance outcomes associated with either full exploration or full exploitation. With either fast- or slow-moving workers, full exploitation is more beneficial than full exploration. This is not surprising. The organization is better off when the manager allows what he or she knows to guide allocation decisions. Acting on beliefs, even if those beliefs are not completely accurate, is better than deciding randomly.

As exploitation increases from its minimum, we see a steady-rise in firm performance. Performance rises more quickly when workers are slow- versus fast-movers. When workers are slow-movers, there is a vertical shift down in performance, when the manager exploits 80% of the firm's financial capital and explores with the remaining 20%. Above 80% exploitation, performance for slow- and fast-moving workers equalizes and remains equal. When the workers are fast-movers, the organization is better off when the manager fully exploits what he or she knows. But when the workers are slow-movers, the organization is better off when the manager exploits a large fraction of the firm's financial capital and uses the remaining amount for exploration. The results in the bottom left panel describe how performance changes when workers in the organization move quickly or slowly and as exploitation varies. The surface plot illustrates the result we have already discussed but illustrate the full range of parameter values. In the lower right corner, we see combinations of mobility and exploitation that result in higher levels of performance.

The results in the right panel of Fig. 2 provide some insight into the performance outcomes we observe. It is no surprise information quality is high when exploration is at its maximum value. When the workers are slow-movers, information quality remains high as exploitation increases. There is a vertical shift down in information quality when the manager exploits more than 80% of the firm's financial capital. The decline in information quality matches the decline in organizational performance. The decline in performance is a byproduct of the manager knowing less and allocating capital poorly. When the workers are fast-movers, as exploitation increases, we see an immediate drop in information quality. Information quality remains low as exploitation continues to increase. We do not see an uptick in performance when information quality is high for fast-moving workers. When workers are fast-movers, the manager must allocate a large proportion of financial capital to exploration before any information benefits are obtained. So much exploration is required to be informed; the manager is unable to capitalize on what he or she learns. When workers are fast-movers, even though there is a decline in information when the manager increases exploitation, we see a steady rise in performance. The rise is lower than the rise in the slow-worker condition. Information quality is lower in the fast-mover condition, but even in this condition, the firm is better off when the manager exploits what he or she knows, even if what the manager knows is of low quality. The manager is better informed when the workers are slow-movers, but even under this condition, the manager is only able to capitalize on what he or she knows, when he or she allows knowledge to guide 80% of the capital allocation decisions and explores with the remaining 20%. The results in the bottom right panel describe how information changes as mobility and exploitation vary. The information surface plot illustrates what we have already discussed but across the full range

of parameter values. In the lower right corner, we see combinations of mobility and exploitation that result in higher information quality. There are more combinations of worker mobility and exploitation that produce high-quality information, but some of those combinations do not allow the manager to utilize what he or she knows. So we do not see an improvement in performance.[5]

Robust networks

The results we have discussed make a critical assumption about the implication of interproject mobility on how members of the firm are connected to each other. For the previous set of results, interproject mobility has a negative effect on network connections. When two individuals stop working together on the same project, their network connection decays quickly. Network connections are fragile. One can imagine there are organizations where mobility will have less of a negative effect on existing connections. We allow network connections to decay more slowly. When network decay is slow, connections are more robust. Even if two individuals are no longer working on the same project, they can still maintain their network connection. When relationships are more robust, a different set of outcomes emerge.

The results are illustrated in Fig. 3. The performance results are in the left panel and the information results are in the right panel. When we focus on performance, again, exploitation is better than exploration. But while there are still differences in firm performance when the organization contains either fast or slow-moving workers, the advantage shifts to the fast-moving worker condition. Moreover, notice the difference in the average and maximum performance when relationships are fragile versus when they are robust. When network connections are robust, the maximum and the average level of performance are higher. As exploitation increases, we see a steady rise in performance in both the fast- and slow-moving worker conditions. Initially, the increase is higher in the fast-moving condition. At 80% exploitation, we see a positive vertical shift in the slow-moving worker condition. After the shift, performance levels remain the same across the two worker conditions. The information graph tells the story. When relationships are robust and the workers are slow-moving, information quality is low with full exploration. Workers do not react to the differences in capital allocation and so the manager does not learn. The manager does not begin to learn until he or she allocates 80% or more of the firm's capital to "high-quality" projects. The more unequal allocation of resources leads to more interproject mobility and more learning. When workers are fast-moving, they are more sensitive to differences in capital allocation and begin to shift projects with small differences in budgets. Information quality increases dramatically with a small increase in exploitation and remains high as exploitation continues to increase. When network connections are robust, it is better for the manager to fully exploit what he or she knows. This is true, independent of the kind of workers

[5]Perhaps it is useful at this point to compare our baseline results to March's findings. Recall March found firm performance (learning in his case) was high when the manager was a fast learner, and the individual workers were slow learners. The slow-learners in March's organization were more autonomous and instead of being influenced by the manager's choices, focused on whatever solution they were developing. The greater autonomy resulted in greater solution space coverage. The dynamics are not the same but the slow-moving workers in our organization accomplish something similar to the slow learning workers in March's framework. Project quality can change. When workers move slowly, the manager and other workers can potentially hear when a project's quality changes. This of course depends on whether there is a network in place to allow for the important information to be communicated to the manager and individual workers in the firm.

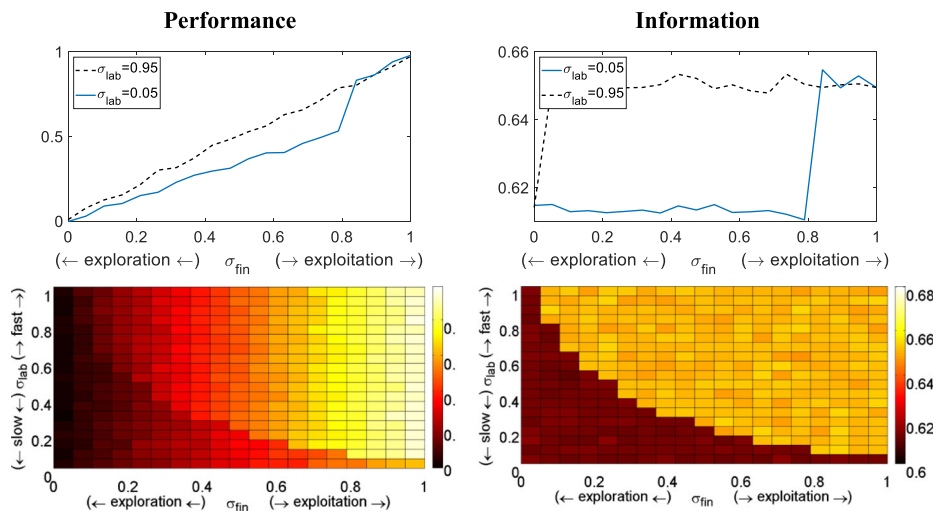

Fig. 3 The figure plots how performance (left) and information (right) vary as a function of the exploration-exploitation balance and labor speed. In the top panels labor is either slow (blue lines) or fast (dashed black lines). Bottom panels plot outcomes for the full parameter space. The difference relative to Fig. 2 is that we are using a calibration with robust networks, instead of fragile

in the organization. This fact is more transparent in the surface plots for performance and information, when network connections are robust. The plots are illustrated in the lower panels for Fig. 3. For every value of worker mobility, we see an increase in performance as exploitation increases. The same pattern holds for information, but the increase is from low to high. When worker sensitivity is high (high mobility), there is an immediate positive vertical shift from low to high information quality. As worker sensitivity declines, the positive shift occurs at higher levels of exploitation. This makes sense. The improvement in information results from workers moving across projects. Workers who are less sensitive need more unequal distributions in capital, which is what exploitation is, before they are willing to move.

Variation in network decay

Figure 4 shows what happens to performance and information if instead of fixing network decay at illustrative values, we allow network decay to vary from 0 (robust

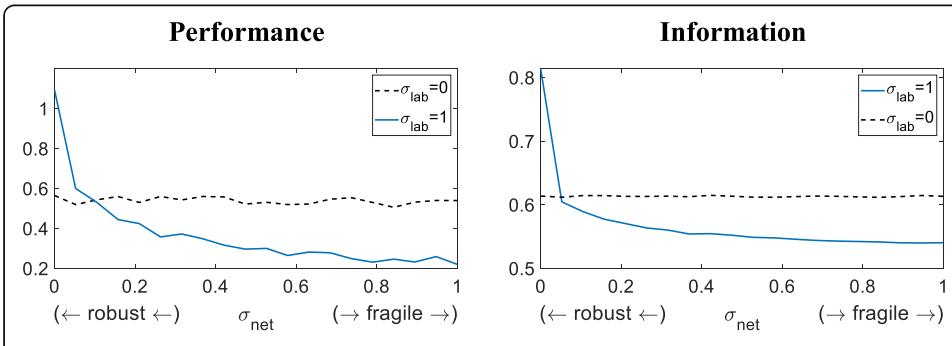

Fig. 4 The figure plots how performance and information vary as a function of network fragility. The exploration-exploitation balance is set at 0.8, and labor is either slow (dashed black lines) or fast (blue lines)

connections) to 1 (fragile connections). Exploitation is set at 80% in all cases. The dashed line is for slow-moving workers and the solid line is for fast-moving workers. The results are quite different depending on how quickly workers move. When workers move slowly, the rate at which network connections decay is irrelevant. Since workers are relatively immobile, network connections remain organized around the projects. For fast-moving workers, variation in network decay matters a great deal. With fast-moving workers, as the network decay rate increases, there is a steady decline in performance. The most dramatic shift occurs with the initial increase in network decay. The results for information quality mirror our performance outcomes. Overall, we see network connections need to be robust enough that the movements between projects allow for individuals to maintain old ties while new ones develop. When this is the case, each worker is likely to be connected directly or indirectly to a large number of projects, and if the manager simply chats with a single worker, the interaction would be highly informative.

Interproject mobility

If network decay is a key contingency, interproject mobility is a key explanatory variable. The rate of network decay shapes how interproject mobility affects the network formation process, which influences information quality and ultimately performance. In this section, we illustrate how. There are two panels in Fig. 5. In the right panel,

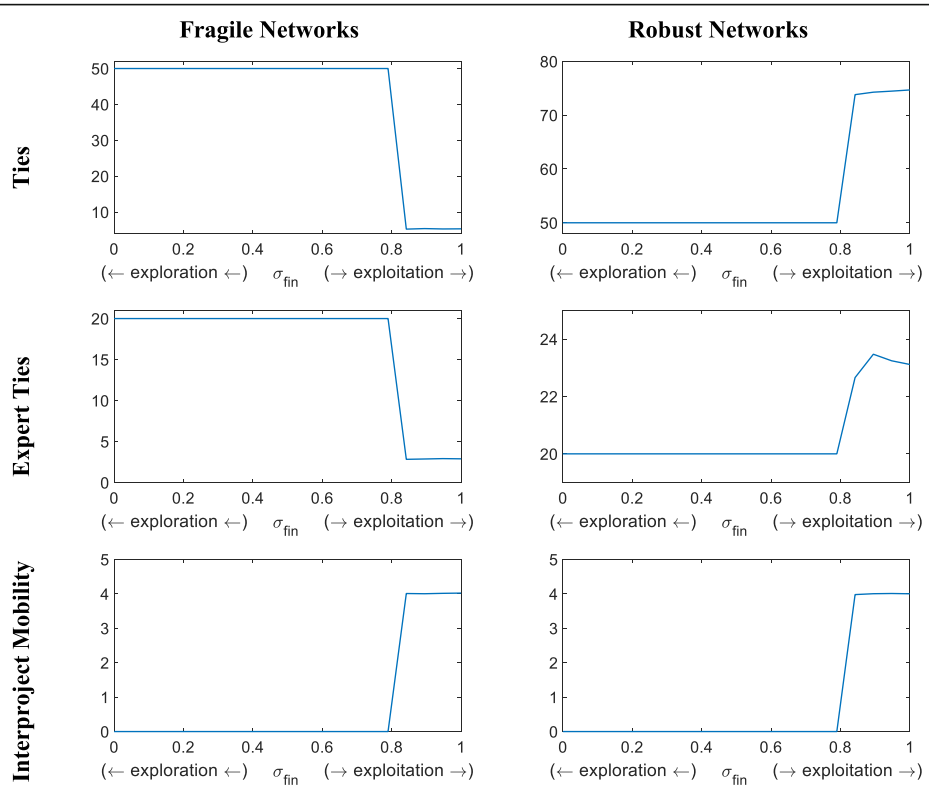

Fig. 5 The figure plots how the exploration-exploitation balance impacts (i) the long-run average for the total number of ties (top row), (ii) the long-run average for the number of ties to experts (middle row), and (iii) the average number of workers switching project per round (bottom row). The left column refers to the fragile networks condition and the right column to the robust networks condition

network connections are fragile; in the left, network connections are robust. In the top and middle rows of both panels are indicators of the organizational network. In the top row is the number of network connections between individual workers. In the middle row, we highlight a special kind of network connection. We show the number of connections between individual workers and experts. The experts are the informational seeds in our analysis and each project has one expert (recall, these workers are assumed to never move). As the number of connections with the experts decline, it is harder for the individual workers to discover the quality of different projects. In the bottom row of each panel is interproject mobility. The variable shows the number of workers who switch projects as exploitation increases. Worker sensitivity is at .05 in both panels.

As the manager shifts from exploration to exploitation, we see little change in the number of switchers. In both panels, individuals stay put until exploitation reaches 80%. At 80%, we see a vertical shift in interproject mobility. While changes in exploitation have the same effect on interproject mobility, mobility has one effect on network connections when ties are fragile and another effect when connections are robust. When network connections are fragile, interproject mobility has a negative effect on the number of connections between individual workers and experts, and also on the number of ties between individual workers. When network connections are robust, interproject mobility has a positive effect on the number of relationships between individual workers and experts, and on the number of ties between the individual workers. When network connections are robust, individuals are able to maintain existing ties and develop new connection when they change projects.

The network connections aggregate to define the broader organizational network. Below 80% exploitation, there is little interproject mobility, and network clusters develop around the projects, as individuals who work on the same project develop relationships with each other. Under both the fragile and robust network connections, the organizational network is composed of network communities and the manager is the bridge between the communities. The organizational network approximates a small world network. The manager can accumulate a considerable amount of information by simply speaking to one individual worker from each network cluster. When network connections are fragile, at levels of exploitation above 80%, we see a dramatic decline in the number of network connections. There is a shift from a small world network to an organizational network that is sparse and disconnected. When network connections are robust, at exploitation levels above 80%, we see a dramatic increase in the number network connections. Since ties decay slowly, as the number of network connections grows, the organizational network goes from a small world to a densely connected network.

The simulation results illustrate the concern which motivated our research project but with an important caveat. We worried that when individuals were goal-oriented, their choices and decisions could have a negative effect on the broader organizational network, potentially creating a tradeoff between exploitation and the ideal network. We were partially correct, and the key contingency is the rate at which existing network connections decay. When network decay is slow, goal-seeking behavior has a positive effect on network formation and ultimately performance. As network decay increases, a tradeoff emerges. The exploitation of knowledge undermines the very network which allowed for the information to be obtained. When network connections are fragile, organizational performance requires a balance between exploitation and exploration.

So far, we have focused on the long run implications of our manager's choices. There is value in examining a specific learning sequence. There are two panels of results in Fig. 6, one when network connections are fragile and another when they are robust. To illustrate the implications of a change in exploitation, we show a manager who exploits 70% of the firm's capital, but on round 500 decides to exploit more than 80% of the firm's capital. A 10-point change in exploitation seems like a small experiment. After the decision, we see a decline in information quality and we also see a corresponding decline in organizational performance. In Fig. 7, we also show how the organizational network reacts to the change. The network goes from being a small world network to a sparse network. When network connections are fragile, the decision to exploit information comes at the expense of the organizational network which allows for superior decisions to be made. The decision to exploit more information results in a vicious cycle, which undermines learning and performance. Current improvements in performance come at the expense of future performance. When network connections are robust, the decision to increase from 70 to 80% exploitation is beneficial. We see an increase in information quality and performance. The organizational network shifts from a small world network to a densely connected network. When ties are robust, the decision to exploit more information creates a virtuous cycle, improving current performance and setting in motion a sequence of events which lead to superior performance in the future.

The manager's network

Our simulation results have illustrated the value of a manager exploiting what he or she knows when network connections are robust and finding a balance between exploitation and exploration when network connections are fragile. We have focused on the network between non-managers. The manager's network has been held small and constant. The network is kept small to illustrate the value relationships between

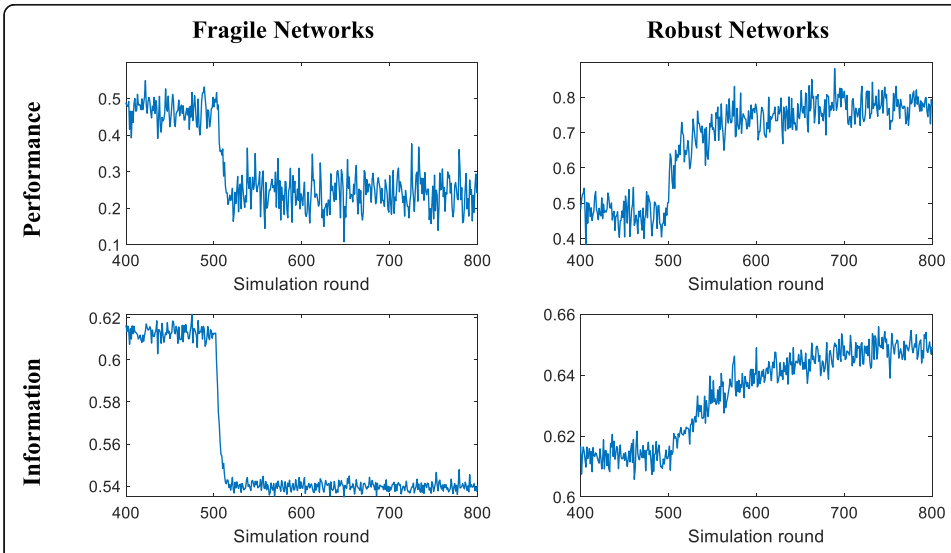

Fig. 6 The figure plots the evolution of performance (top panels) and information (bottom panels) for the two network conditions, fragile (left panels) and robust (right panels). The exploration-exploitation balance is set at 0.7 before round 500 and at 0.8 afterwards

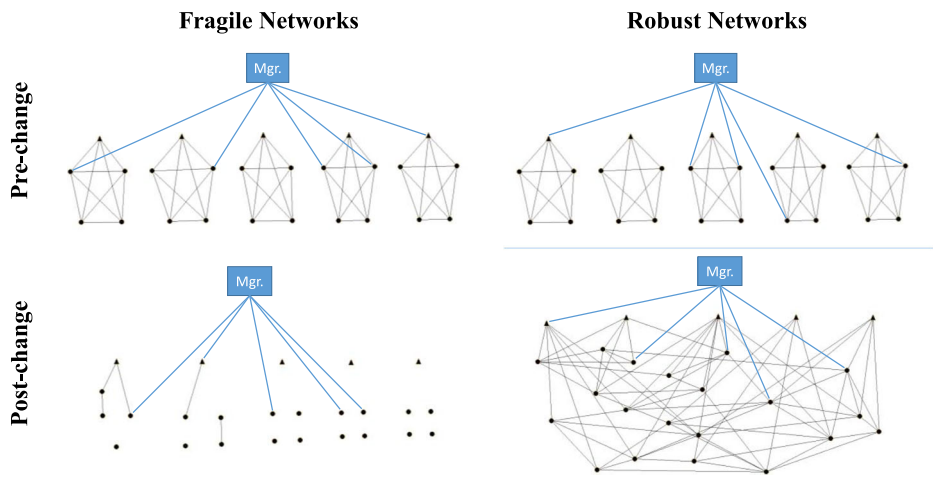

Fig. 7 The figure plots illustrative social networks before (top panels) and after (bottom panels) the change in the exploration-exploitation balance from 0.7 to 0.8, under the fragile networks condition (left panels) and robust networks condition (right panels). The rectangle represents the manager, the triangles represent the experts, and the circles represent the non-expert workers

workers can create. This is not to suggest the manager could not benefit from communicating with more workers. The value of the manager communicating with more workers is especially evident when the organization is large.

Figure 8 illustrates how firm performance and information varies as the manager interacts with a larger fraction of the organization. There are two lines in the figure. The one depicted in blue has the same size as the organizations depicted in previous figures

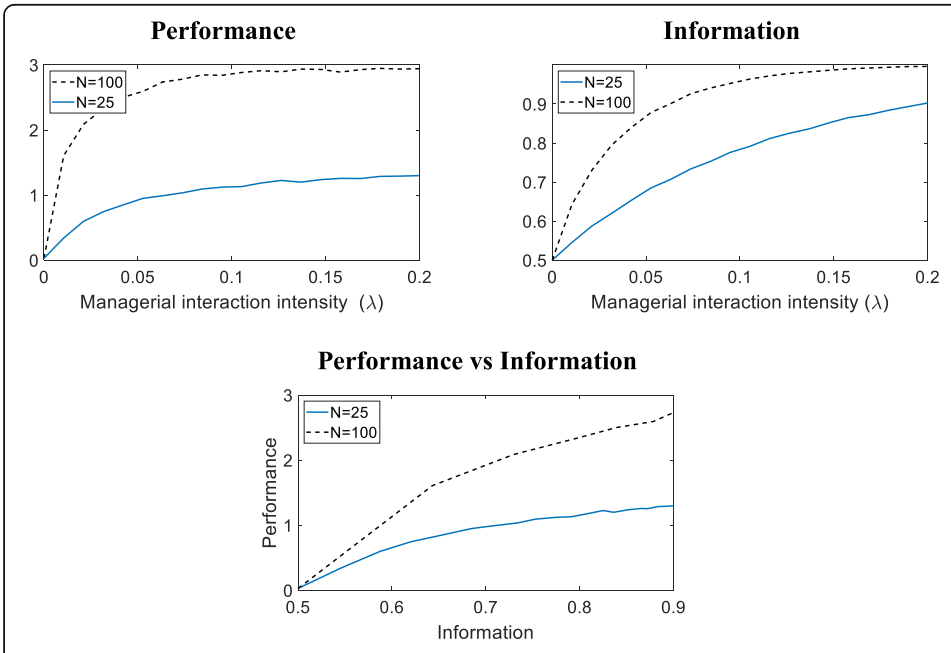

Fig. 8 The top panels plot performance (left) and information (right) for varying levels of managerial interaction intensity, and for two types of organizations, large (dashed black line) and small (blue line). The bottom panel combines the top two panels, and represents information-performance pairs for each organization type.

(25 workers). In contrast, the dashed lines represent a much larger organization (100 workers). The important features of the organization remain the same, equal number of projects, experts, and the same overall amount of financial capital. Performance and information both increase as the manager interacts with a larger fraction of the firm. Performance increases faster than information. The results indicate that information is more valuable in the larger organization. In the larger organization, the returns to the same amount of information are much higher. For example, when the manager communicates with 10% of the firm, there is a 20% improvement in information quality, but there is approximately a 200% increase in performance. The point is more transparent in the bottom panel, where we show the association between information and performance for the large and small organization. A better-informed manager makes financial capital allocation decisions that lead to more effective matching of workers and projects. More effective matching pays off more in the larger organization because there are more resources to match. This is a somewhat obvious point but emphasizes the kind of scale economies an informed manager can create. We make this point because the same is true for dynamics that produce a better organizational network.

Learning to balance

The manager we are analyzing actually faces multiple learning challenges. The first challenge is to identify high-quality projects. Our manager uses his or her interactions with individuals to develop more accurate beliefs about project quality. The more workers are informed and the more interactions the manager has with the workers, the more accurate are the manager's beliefs. More accurate beliefs equal higher quality information, and with everything else constant, higher quality information improves performance.

The second learning challenge is that the manager needs to realize that there exists an optimal exploration-exploitation balance, which is contingent on the organizational context, in particular network fragility. In other words, the manager needs to be aware that trade-offs between exploration and exploitation exist: exploit too much and, at least in some contexts (fragile relationships), you end up undermining the network that is the foundation for being informed about project quality to begin with.

Even if our manager is able to learn project quality and understands the potential for the exploitation-network tradeoff, the manager must find the optimal balance between exploration and exploitation in his or her specific context. We focus on how our manager can learn how to address these challenges. We assume our manager attempts to learn by experimenting with different levels of exploitation.

The more difficult task appears to be learning the appropriate amount of exploitation when network connections are fragile. We consider two ways the manager could attempt to learn. The manager could experiment with exploitation and focus on how changes in exploitation affect performance, or the manager could focus on how changes in exploitation affect information. It turns out that it is easier for a manager to learn from experimenting with exploitation, if he or she focuses on information instead of performance.

Consider Fig. 6, where our manager experiments at round 500 with an increase from 70 to 80% exploitation. Focus on the left panel, the fragile network condition. The

increase from 70 to 80 crosses a critical threshold, which leads to too much interproject mobility, ultimately undermining the network foundation for managerial information about project quality. The question is, could our manager quickly learn the increase was a mistake, by observing the change in performance? The answer is no, and the reason is illustrated in the left panel of Fig. 9. Short-run performance is too noisy, the standard deviation being orders of magnitude of the mean.

To make the point clear, let us compute a rough estimate of how many rounds it would take the manager to learn that the increase in exploitation was a mistake. Average performance drops by about 0.25, and the standard deviation of performance is about 1.1 if we take the average standard deviation before and after the change. For simplicity and ease of exposition, let us assume this is the standard deviation the manager will work with when attempting to statistically assess whether performance changes were significantly different from zero. The number of rounds that would be required for a t test of 2 would be the one that produces a standard error of 0.125 (0.25/2). With a standard deviation of 1.1, this is about 70 rounds:

$$1.1/\sqrt{70} \approx 0.125$$

If 70 rounds is too large or too small depends on the time scale one has in mind. With no specific time scale in mind, we ask if there is a better way to learn. Let us now switch to analyzing what happens to information quality, i.e., the accuracy of our manager's beliefs. Assume the manager has a way to compare his or her beliefs with the truth from previous rounds. If we do the same exercise we did with performance, the change in mean managerial information at round 500 is about 0.08 (0.62–0.54) (right panel of Fig. 9, blue line). The standard deviation is about 0.075, again averaging across before and after the change. This means that in order to get a standard error of about 0.04 (0.08/2), you need only about 3 rounds:

$$0.075/\sqrt{3} \approx 0.04$$

This stands in stark contrast with the 70 rounds that it takes for the manager to learn about performance. The key implication is that for a manager who wants to learn about the optimal level of exploitation, the task is much more easily done by looking for differences in information quality. This assumes that the manager is able to know, even if

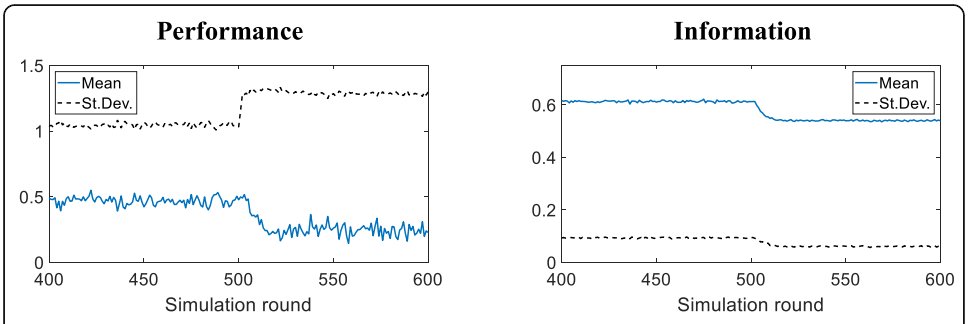

Fig. 9 The figure plots round-level means and standard deviations for performance (left panel) and information (right panel). The exploration-exploitation balance is set at 0.7 before round 500 and at 0.8 afterwards

only later on, if her beliefs about a particular project were right or not (otherwise she would not be able to compute the accuracy variable plotted in the figure); but at least in some settings, we believe this is a plausible assumption. We have focused on learning in the fragile network condition because it represents a more difficult learning environment. Our manager could use the same approach to learn the optimal amount of exploitation in the robust network condition. Our previous analysis shows that if the manager starts with a focus on information, he or she should be able to use selective trial and error experimentation to distinguish similar levels of information and to eventually discover the appropriate amount of exploitation.

Summary and discussion

Our research was motivated by a basic observation. Would individual choices undermine the benefits learning in an organizational network can create? Prior research on mutual learning in an organizational network has emphasized the importance of a balance between fast and slow learning processes. For March, the contrast was between a manager and individual workers. Superior outcomes were obtained when the manager was a fast learner and the workers were slow learners. In more recent work, everyone learns at the same rate, and the balance occurs between the kind of problem individuals are solving and the rate at which knowledge and information spreads among them. Some networks allow knowledge and information to spread quickly while other networks allow for knowledge and information to spread more slowly. When individuals are solving a basic problem, there is no need for balance. Learning is a walk uphill and an organizational network that allows for fast transmission increases the speed of the walk. When individuals are solving a complex problem, superior outcomes are obtained in an organizational network that allows for some knowledge transfer but also places limits. Superior learning occurs when there are buffers between different parts of the organization, with fast learning within but slow learning between the different parts. An example of an ideal network is a small world network. More recent research has allowed for networks to be endogenous. While we see significant merit in this line of research, we worry it has neglected the potential for a bad outcome to be obtained. With choice, it is possible individuals could move out of networks which allowed for the superior learning outcomes to be obtained. The point is if individuals are goal-oriented, the choices they make could result in them selecting out of a network which enhances learning and into a network which undermines learning.

Our simulation results illustrate that there is in fact a potential for the tradeoff we imagined. But we are only partially correct. The key contingency is network fragility. When network connections are robust, we did not find a tradeoff between exploitation and the organizational network. Indeed, we found a virtuous cycle. Exploitation leads to the formation of a dense network, improving the quality of the information available to the manager, leading to better quality decisions and higher organizational performance. When network connections were fragile, exploitation produced a sparse disconnected network, reducing information quality, and ultimately performance. When network connections were fragile, a manager had to find the right balance of exploitation and exploration. The manager must find a balance which limits the rate at which workers move from one project to the next, allowing the manager to exploit some of what he or she knows, without undermining the very network which allows for useful

information to be obtained. Our results also illustrate an important dynamic between exploitation and exploration at two levels of analysis. When network connections are robust, our manager can fully exploit what he or she knows. When the manager exploits what he or she knows, the decision encourages our workers to leave their current projects and move to different projects. Exploitation by our manager leads to exploration by our individual workers.

Our simulation results have important implications for scholars with an interest in mutual learning in networks. Our findings illustrate when allowing for an endogenous network process can become problematic. Note, however, we are not suggesting that scholars should follow our approach. We have focused on network formation as a byproduct of purposive action. In our framework, network connections emerge and decay as a byproduct of individuals attempting to find better places to work in our organization. The individuals are not looking for better network connections. Anjos and Reagans (2013, p. 12) have shown the potential for a tradeoff between exploitation and exploration when individuals are looking for better network connections. In their framework, the extent of the tradeoff depends on if individuals act on what they know quickly or at a more moderate pace. We see significant merit in frameworks and models that allow for goal-oriented individuals to influence network formation, learning and ultimately performance (Clement and Puranam 2018; Songhori and García-Díaz 2018).

Our simulations also have implications for scholars with an interest in organizational learning defined more broadly. To illustrate how, we conclude with a brief discussion of the Carnegie School, which has heavily influenced organizational learning research. The Carnegie School was characterized by two alternative views of the organization. In the view articulated by March and Simon (1958), an organization could be designed to achieve a specific goal. March and Simon highlighted the importance of developing organizational routines (standard operating procedures) that "bounded" the individual decision-making process, aligning the outcomes of those decisions with the broader organizational (or subunit) goal and objective. Organizational learning was achieved as individuals become more efficient and effective decision-makers. Cyert and March (1963) articulated an alternative view of the organization. In their view, the organization was a political machine. Instead of one goal, there were many goals in an organization, often in conflict with each other. The organization was a political coalition that was characterized by a decision-making process that allowed individuals with conflicting goals and objectives to discover which activities were worth pursuing. The goal was not fixed. It was endogenous. Slack resources were a critical part of the process.

With hindsight, one could say that one view focused on how an organization could exploit what it knew, while the other view focused on an organization designed to explore and discover new activities. Our model combines elements of both views. The individuals in our organization are characterized by bounded rationality. They are goal maximizers but they use simple heuristics to maximize their goals. Our manager follows a very simple decision-making heuristic. Our workers use a basic heuristic for deciding where to work. While we have focused on exploitation and exploration in our discussion, we measure exploitation with an indicator of financial slack (Rajan et al. 2000). Our model provides a point of integration between the two views of organizational learning from the Carnegie School. Our manager decides how much of

the firm's capital should be exploited and how much should be used as slack. The decision sets in motion a chain of events which affect network formation, learning, and ultimately performance. In the right context, our boundedly rational decision makers could discover and re-discover which projects were worthwhile. They could exploit what they knew, while continuing to search for more attractive opportunities.

Appendix: The formal model

The objective of this section is twofold. First to offer a formal detailed description of the model, in order to be as transparent as possible about our assumptions. Second, to discuss these assumptions in light of our main results. Before jumping into the description, it is useful to group the parameters that govern our simulations into two categories: fundamental and auxiliary. Fundamental parameters are those that are critical for our arguments, for example the parameters that describe exploitation and worker mobility. Auxiliary parameters are those which we believe we should be explicit about, not only to keep the model more general, which is useful for subsequent research; but also because it facilitates presentation. This categorization of parameters notwithstanding, some auxiliary parameters are actually critical for our results, however discussing this in the main text would unnecessarily complicate our discussion. Such discussion is therefore one of the purposes of this section, and it will help further delineate our scope conditions.

The Organization, its actors, and the performance function

In our agent-based model there are M projects, indexed by m; and N workers, indexed by i, each having a default assignment; and one manager who controls how financial capital is allocated. All projects are symmetric, in the sense that a similar number of workers are assigned to each project initially; and the statistical properties of project quality are the same for all m. As described in the main text, most of our analyses refer to an organization where N=25 and M=5.

Total performance: The total performance of our organization at each point in time, denoted as Q_t, is the sum of project-level performances, $Q_t = \Sigma_m q_{m,\ t}$; where $q_{m,\ t}$ equals the performance of project m at time t:

$$q_{m,t} = k_{m,t} \times v_{m,t} \times \left[1 + \phi \times \sum_{i=1}^{N} a(i,m)\right] \tag{1}$$

$k_{m,\ t}$ is the amount of capital the manager invests in project m for time period t. The investment decision is made at the end of time period *t-1*. Our manager has a fixed budget, which we set equal to 1,[6] and invests the full amount of capital; so $\Sigma_m k_{m,\ t} = 1$.

$v_{m,\ t}$ refers to the quality of project m at time period t. In our organization, at any point in time t, the quality of project m can take on one of two values, a low value v_L or a high value v_H. To capture the idea of investment under uncertainty, we allow the quality of project m to change. The probability of project m remaining at the same level of quality during the time interval *t-1* and t is given by $\gamma_v \in [0.5,1]$. All of our projects

[6]Our results are however independent of the total size of the budget, given the linearity of the performance equation in capital allocations.

are parameterized in the same way and the changes in their quality are statistically independent from one another. Technically, project quality follows a two-state Markov chain (for a similar model, see Anjos and Reagans 2013). To illustrate how this process works, consider a simple example where the organization has two projects and $\gamma_v = 0.7$. Further assume that at time t both projects are in the high-economic-potential regime, meaning $v_{1,\ t} = v_{2,\ t} = v_H$. Then in time period $t+1$ both projects remain in the same regime with probability 0.7^2; both projects become low-potential with probability 0.3^2; and only one of the projects changes potential with probability $1\text{-}0.7^2\text{-}0.3^2 = 0.42$.

The closer γ_v is to 1, the more stable a project's quality.[7] In our examples we always employ a calibration where $\gamma_v = 0.7$, but our results would go through as long as γ_v is distant enough from both 0.5 and 1. If γ_v is too close to 1, there is no uncertainty and learning is trivial (done after a few rounds); on the other hand if γ_v is too close to 0.5 there is too much uncertainty and knowledge becomes stale so quickly that it is hard for the manager to use it in a timely and useful way.

In our calibration we set $v_L = -1$ and $v_H = 1$. This means that without any information, the long-run average performance of the organization would be 0 (see equation (1) and set the average $v_{m,\ t}$ at 0). This is mostly a useful normalization, although the wedge between v_L and v_H does affect results. In particular this wedge influences both the volatility that the organization is exposed to, as well as the returns to information (consider the extreme case where $v_L = v_H$; then there is no need to learn).

Finally, the last term in equation (1) captures labor's contribution to performance. The term $a(i, m)$ is an indicator function taking the value of 1 if worker i is working in project m (and 0 otherwise), and ϕ scales the contribution of labor to performance. In our calibration we set this parameter to 0.1, however our results obtain for a range of values. The role played by ϕ is as follows. As ϕ becomes close to zero, the only channel through which worker movements affect performance is through learning. In this extreme case, worker movement has no direct impact on performance, yet it still affects whether and how networks form. As ϕ increases, then worker movements also impact performance directly; the more workers a project has, the higher its performance, as shown in equation (1). Moreover, if the manager is well-informed and frequently allocates much financial capital to high-quality projects, then labor amplifies the performance returns to managerial knowledge via workers flocking precisely to these projects. In other words, a better informed manager also fosters a better matching of workers and projects.

Manager beliefs and capital allocation: At the beginning of every time period t, the manager holds the belief $b_{m,\ t}$ about the quality of project m. This belief is the subjective probability that project m has high quality. As the manager gathers information, the subjective belief is updated. We describe this process in greater detail in the Information Environment section below. Currently, we are concerned with how the manager's subjective belief about a project's quality translates into an amount of allocated capital. Our manager uses a simple heuristic:

[7]The likelihood that a project changes from low to high is governed by the same parameter, which immediately implies that the unconditional (or long-run) probability of either state is 1/2.

$$k_{m,t} = \begin{cases} \left(1 - \sigma_{fin}\right) \times \dfrac{1}{M} & if \;\; b_{m,t} \neq \max\limits_{j}\left\{b_{j,t}\right\} \\ 1 - (M-1) \times \left(1 - \sigma_{fin}\right) \times \dfrac{1}{M} & if \;\; b_{m,t} = \max\limits_{j}\left\{b_{j,t}\right\} \end{cases} \tag{2}$$

where parameter $\sigma_{fin} \in [0, 1]$ captures the manager's tendency to exploit what he or she knows, which we interpret as a stable feature of the organization. The project with the maximum subjective belief is always allocated the most financial capital. How much more than the remaining projects depends on σ_{fin}. If σ_{fin} is 1, there is a "tight" link between allocated resources and quality. This approach corresponds to an exploitation strategy, where the manager maximizes the immediate payoff of his or her current knowledge. If σ_{fin} is 0, capital is allocated across projects equally, independent of quality. This approach corresponds to complete exploration.

Worker mobility and network formation

Worker mobility: We assume that there is one expert per project, and experts always remain in their default assignment. In contrast, individual workers can move from their default project to one they prefer more. In particular, the individual workers in our organization prefer to work on projects which have been allocated a larger amount of capital. The allocated amounts of capital are common knowledge in our organization. The details of our modeling in this regard are as follows. The worker starts in his or her default assignment and decides if he or she would prefer to work on a different project. Specifically, at every time period t each worker is presented with an alternative project assignment (randomly drawn from the other M-1 projects); and in case the budget of the default assignment is smaller than σ_{lab} times the budget of the alternative assignment, then the worker switches (for that time period). Parameter $\sigma_{lab} \in [0, 1]$ therefore measures how strongly workers prefer to work on projects with more funding. If $\sigma_{lab} = 0$, budget-size motives are completely turned off (slowest possible labor), if $\sigma_{lab} = 1$ then the smallest difference in budget will make the focal worker switch projects (fastest possible labor).

Network formation and decay: We model network formation and decay as a function of the amount of time individuals spend working on the same projects. Reagans and McEvily (2003, pp. 252–253) describe the positive effect working on the same projects can have on network formation. Time spent working together is an organizational equivalent of physical proximity. Proximate individuals have more opportunities to develop and maintain a network connection. Our network connections form and decay as time spent working together changes. At the beginning of time period t, if two individuals were connected at *t-1 and* continue to work on the same project, their network connection is maintained. If two individuals were not connected at *t-1*, but are working on the same project at time period t, a tie will form between the two individuals with probability $\alpha \times overlap$, where "overlap" measures the number of consecutive time periods the workers have been together on the same project.[8] For example, suppose α = 0.2 and two workers who join the same assignment are disconnected. After the first time period, there is a 20% probability a network connection will develop between them. If the tie is not formed at the end of the first time period together, then there is

[8]In rigor, the likelihood of the tie forming is $\min\{1, \alpha \times overlap\}$, since probabilities need to lie in the unit interval.

a 40% probability a network connection will develop at the end of the second time period, and so on. In our calibration we set α = 0.01. This parameter naturally plays an important role in the model. If α is too close to 1, then ties form immediately and in a sense it does not matter much how workers move; a minimal amount of overlap makes the learning network emerge. On the other hand, if α is too close to zero, then the social network is always poor (the initial network is completely disconnected).

If two individuals who were connected at *t-1* are no longer working on the same project at time period *t*, their network connection will decay with hazard rate σ_{net}. This is the parameter which in the main text we referred to as network fragility, since a higher σ_{net} makes network connections more responsive to worker movement.

The information environment

Learning: The mechanisms through which information is disseminated and eventually affects managerial allocations are summarized as follows.

At the beginning of every time period, the manager gathers information from a worker i about project m with probability λ. In our main calibration this parameter is set to 0.05. For the results to be interesting, λ cannot be too close to zero (no managerial knowledge possible); neither too close to 1 (all managerial knowledge is obtained directly from experts, and there is no role for the social network).

We assume that both experts and generalist workers truthfully convey their beliefs to the manager. Given a set of worker beliefs, the manager adopts the most informative belief, which in our case is the most extreme belief. This amounts to rational decision-making in our case.

At every time period, both experts and individual workers also collect their contacts' opinions about all projects; and adopt the most informative beliefs about a project (potentially their own). If an actor does not obtain new relevant information about project m at a given round, his or her beliefs about its quality mean-revert in accordance with the true stochastic process governing $v_{m,\ t}$ (assumed to be common knowledge in the organization).

Belief dynamics: Let us denote by $b_{i,\ m,\ t}$ the subjective probability that actor i (expert or individual worker) assigns to project m being high-potential (i.e., that $v_{m,\ t} = v_H$). Our mechanism of information transmission is as follows:

1. Experts are always informed about the true $v_{m,\ t}$ of their own project. This is the informational seed.
2. Given the social network at time period t, every actor communicates his or her existing beliefs $b_{i,\ m,\ t}$ to all his or her contacts. This is done before updating beliefs with others' opinions. A way to conceive of the mechanism is that everyone first writes and sends out all "recommendation letters"; and everyone opens recommendation letters only at the end. In this way, there is no need to specify the timeline of communication.
3. Given a set of beliefs received by i from his or her contacts, i adopts the one that is the most extreme, i.e., farthest away from 1/2 (the unconditional probability that a project is high-potential). This makes sense in our setting where beliefs are taken to be rational; one can think about extreme beliefs intuitively as the beliefs of those actors who have seen the true state more recently.

4. If no information is transmitted during time period t, then beliefs evolve according to the following law of motion:

$$b_{i,m,t} = \delta b_{i,m,t-1} + (1-\delta)\frac{1}{2} \tag{3}$$

where

$$\delta = 2\gamma_v - 1. \tag{4}$$

δ gages how much weight actor i places on past information vs. simply taking an agnostic/unconditional view (probability of 1/2). Note that equation (3) implies that as time goes by the subjective beliefs held by i about m converge gradually to 1/2 (assuming no information is observed in the meantime). A "rational" (or Bayesian) actor would want to update beliefs according to (3), since that is consistent with our probabilistic assumptions: Without the arrival of new information, it needs to be the case that

$$\Pr\{v_{m,t} = v_H\} = \underbrace{\Pr\{v_{m,t-1} = v_H\}\gamma_v}_{\text{remains high}} + \underbrace{\left(1 - \Pr\{v_{m,t-1} = v_H\}\right)\left(1-\gamma_v\right)}_{\text{becomes high}} =$$

$$= \Pr\{v_{m,t-1} = v_H\}\underbrace{\left(2\gamma_v - 1\right)}_{=\delta} + \underbrace{1-\gamma_v}_{=(1-\delta)/2}.$$

Our assumptions regarding belief dynamics, albeit in the spirit of the rational actor framework, were not dictated by a strict adherence thereto. Instead, we adopted a very simplistic assumption regarding how individuals construct stocks of subjective knowledge; and this simplistic assumption facilitates our analysis. In further work it would certainly be interesting and relevant to understand how the various organizational dynamics interact with alternative mechanisms of belief formation, with potential implications for learning and performance outcomes.

Acknowledgements
The authors would like to thank the editor and referees at the Journal of Organizational Design for various insights that importantly improved our paper. We are also thankful for the feedback we received on an earlier version of our work from participants in the economic sociology workgroup at MIT (Sloan), participants at the Carnegie Mellon (Tepper) OBT seminar, and participants at the conference honoring Jim March (Carnegie Mellon). Finally, we are indebted to Bob Gibbons for his encouragement, guidance, and patience.

Authors' contributions
Both authors contributed to the analysis and text contained in the manuscript. The author(s) read and approved the final manuscript.

Funding
This work was funded by Fundação para a Ciência e a Tecnologia (UID/ECO/00124/2013, UID/ECO/00124/2019 and Social Sciences DataLab, LISBOA-01-0145-FEDER-022209), POR Lisboa (LISBOA-01-0145-FEDER-007722, LISBOA-01-0145-FEDER-022209) and POR Norte (LISBOA-01-0145-FEDER-022209).

Availability of data and materials
Data sharing not applicable to this article as no datasets were generated or analyzed during the current study.

Competing interests
The authors declare that they have no competing interests

Author details
[1]NOVA School of Business and Economics, Campus de Carcavelos, Rua da Holanda, n.1, 2775-405, Carcavelos, Portugal.
[2]MIT Sloan School of Management, 100 Main Street, Cambridge, MA 02142, USA.

Received: 26 November 2019 Accepted: 25 August 2020
Published online: 06 October 2020

References

Anjos F, Reagans R (2013) Commitment, learning, and alliance performance: a formal analysis using an agent-based network formation model. J Math Sociol 37(1):1–23
Argote L, Ingram P (2000) Knowledge transfer: a basis for competitive advantage in firms. Organ Behav Hum Decis Process 82(1):150–169
Boudreau KJ, Lakhani KR (2015) "Open" disclosure of innovations, incentives and follow-on reuse: theory on processes of cumulative innovation and a field experiment in computational biology. Res Policy 44(1):4–19
Clement J, Puranam P (2018) Searching for structure: formal organization design as a guide to network evolution. Manag Sci 64(8):3879–3895
Cowan R, Jonard N (2004) Network structure and the diffusion of knowledge. J Econ Dyn Control 28(8):1557–1575
Csaszar FA, Siggelkow N (2010) How much to copy? Determinants of effective imitation breadth. Organ Sci 21(3):661–676
Cyert RM, March JG (1963) A behavioral theory of the firm, Englewood Cliffs
Denrell J, Fang C, Levinthal DA (2004) From T-mazes to labyrinths: learning from model-based feedback. Manag Sci 50(10): 1366–1378
Diwas K, Staats BR, Gino F (2013) Learning from my success and from others' failure: evidence from minimally invasive cardiac surgery. Manag Sci 59(11):2435–2449
Epple D, Argote L, Devadas R (1991) Organizational learning curves: a method for investigating intra-plant transfer of knowledge acquired through learning by doing. Organ Sci 2(1):58–70
Epple D, Argote L, Murphy K (1996) An empirical investigation of the microstructure of knowledge acquisition and transfer through learning by doing. Oper Res 44(1):77–86
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
Fu CJ (2020) Draw near to go far: the role of convergence in capitalizing on exploration. Working Paper, MIT Sloan
Golub B, Jackson MO (2010) Naive learning in social networks and the wisdom of crowds. Am Econ J Microeconomics 2(1):112–149
Hasan S, Koning R (2020) Designing social networks: joint tasks and the formation and endurance of network ties. J Organ Design 9(1):1–19
Ingram P, Simons T (2002) The transfer of experience in groups of organizations: implications for performance and competition. Manag Sci 48(12):1517–1533
Kauffman SA (1993) The origins of order: self-organization and selection in evolution. Oxford University Press, USA
Knudsen T, Srikanth K (2014) Coordinated exploration: organizing joint search by multiple specialists to overcome mutual confusion and joint myopia. Adm Sci Q 59(3):409–441
Lazer D, Friedman A (2007) The network structure of exploration and exploitation. Adm Sci Q 52(4):667–694
Levinthal DA (1997) Adaptation on rugged landscapes. Manag Sci 43(7):934–950
Levinthal DA, March JG (1993) The myopia of learning. Strateg Manag J 14(S2):95–112
March JG (1991) Exploration and exploitation in organizational learning. Organ Sci 2(1):71–87
March JG, Simon HA (1958) Organizations
Mason W, Watts DJ (2012) Collaborative learning in networks. Proc Natl Acad Sci 109(3):764–769
Puranam P, Raveendran M, Knudsen T (2012) Organization design: the epistemic interdependence perspective. Acad Manag Rev 37(3):419–440
Rahmandad H (2019) Interdependence, complementarity, and ruggedness of performance landscapes. Strategy Science 4(3): 234–249
Rajan R, Servaes H, Zingales L (2000) The cost of diversity: the diversification discount and inefficient investment. J Financ 55(1):35–80
Reagans R, McEvily B (2003) Network structure and knowledge transfer: the effects of cohesion and range. Adm Sci Q 48(2): 240–267
Reagans, R., & Zuckerman, E. (2008). A shortcut to efficiency: implications of a small but stratified world. Working Paper, MIT Sloan.
Schilling MA, Fang C (2014) When hubs forget, lie, and play favorites: interpersonal network structure, information distortion, and organizational learning. Strateg Manag J 35(7):974–994
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
Simon H (1962) The architecture of complexity. Proc Am Philos Soc 106(6):467–482
Small ML, Adler L (2019) The role of space in the formation of social ties. Annu Rev Sociol 45:111–132
Songhori, M. J., & García-Díaz, C. (2018). Collective problem-solving in evolving networks: An agent-based model. In 2018 Winter Simulation Conference (WSC) (pp. 965-976). IEEE.
Sutton RS, Barto AG (1998) Reinforcement learning: an introduction (Vol. 1, no. 1). MIT press, Cambridge
Watts D, Strogatz S (1998) Collective dynamics of 'small-world' networks. Nature 393:440–442
Weinberger E (1990) Correlated and uncorrelated fitness landscapes and how to tell the difference. Biol Cybern 63(5):325–336
Zuckerman EW, Sgourev SV (2006) Peer capitalism: parallel relationships in the US economy. Am J Sociol 111(5):1327–1366

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.