AARHUS UNIVERSITY

The Ikea-Effect in Collective Problem-Solving: When Individuals Prioritize their Own Solutions

Abstract

To improve problem-solving performance, individuals can rely on social learning. This approach is constrained by an individual's social network which influences the efficiency of the problem-solving process. To date, research disagrees on what kind of network structure is preferable, providing support for efficient network structures, as well as for inefficient networks. However, studies implicitly assume that solvers always imitate superior solutions, an assumption that lacks empirical grounding. We propose a simple derivation of an existing simulation framework by acknowledging a known cognitive bias ('Ikea-effect'), by which individuals are assumed to prioritize individual information. This effect allows inefficiencies to be embodied at the individual micro-level, which reduces the need for inefficiencies at the structural macro-level. Simulation results explain discrepancies in previous results, illustrating how more realistic micro-level assumptions substantially impact macro-level outcomes.

Keywords: behavioral strategy; NK models; social networks; agent-based model; exploration and exploitation

Introduction

Scientific optimization challenges such as folding proteins (Cooper et al., 2010), quantum computing (Heck et al., 2018) and innovation contests (Adamczyk, Bullinger, and Moslein, 2012; Boudreau, Lacetera, and Lakhani, 2011) are examples of highly complex problems that cannot be solved analytically. Due to their scope, such challenges are often tackled by a collective of solvers, e.g. an open crowdsourcing contest or teams within or across organizations. When solving these problems, individuals balance individual and social learning. Individual problem-solvers can engage in independent search for a solution through trial-and-error experimentation that provides feedback and learning (Boudreau and Lakhani 2015). In contrast, social learning allows the opportunity to observe, imitate and learn from other problem-solvers that also search for solutions to the same problem (Mason and Watts 2012; Boyd et al. 2011). *Collective* problem-solving refers to 'the pooling of individual abilities to solve a problem' (Lazer and Friedman, 2007, p. 667), an approach widely used in innovation management, with examples ranging from New Product Development teams to online crowdsourcing initiatives, whether internal or external to the firm. In line with Nickerson and Zenger's problem-solving perspective of the firm (2004) we focus on how, in the absence of hierarchies, such a collective of solvers should be organized; e.g., should you be able to see the solutions developed by others?

Both experimental work and simulation studies have demonstrated how strategically varying the access to solutions in the collective through structural properties of the network influences search dynamics and problem-solving performance (Barkoczi and Galesic, 2016; Derex and Boyd, 2016; Fang, Lee, and Schilling, 2010; Lazer and Friedman, 2007; Mason and Posen et al. 2020; Watts, 2012; Schilling and Fang, 2014). Yet, it is still not clear if a manager should prioritize efficient networks that facilitate quick diffusion of good solutions (Mason and Watts, 2012) or maintain diversity in the system by prioritizing inefficient networks that limit access to the solutions of others' (Bernstein, Shore, and Lazer, 2018; Derex and Boyd, 2016; Girotra, Terwiesch, and Ulrich, 2010; Lazer and Friedman, 2007; Schilling and Fang, 2014). Not only is the evidence mixed, research so far has struggled to isolate and identify the underlying theoretical mechanisms driving the long-term behavioral dynamics. We add two new perspectives to the discussion in order to explain the current discrepancy of results.

First, we question a core assumption of the aforementioned simulation studies of complex problem-solving. In these studies, agents always copy other agents' superior solutions, implying a rationalized assumption. However, empirical studies in a range of fields highlight that this assumption may not be realistic since individuals have a tendency to prioritize their own information over social information, thus not indiscriminately copying any superior solution (Eriksson and Strimling, 2009; Hargadon and Bechky, 2006; Mason and Watts, 2012; Acerbi et al., 2016; Muthukrishna et al. 2016). Second, simulation and experimental studies have focused on improving the average performance of the collective (e.g. Lazer and Friedman 2007; Mason and Watts, 2012; Posen et al. 2020). This focus makes sense when the emphasis is, for instance, the optimization of daily operations. Yet, in the innovation management literature (e.g. innovation contests) one would prioritize how to find novel superior solutions, where it is of less importance if *one* or *all* members of the collective have found the global peak (Boudreau et al., 2011; Girotra et al., 2010). To our knowledge, we are the first to explicitly compare maximum vs. average outcomes in a simulation study.

Based on the widely used NK setup (Levinthal, 1997), the aim of this paper is to examine how different assumptions about individual search behavior and choice of outcome measurements influence collective problem-solving performance. In other words, how should a manager organize a collective problem-solving effort, such as an innovation contest for example? Using an agent-based model, we compare efficient and inefficient networks under three plausible and widely used behavioral assumptions: 1) copy the best (Lazer and Friedman, 2007; Schilling and Fang, 2014), 2) copy the majority (Barkoczi and Galesic, 2016; Fang et *al.,* 2010) and 3) prioritize one's own solution. The latter strategy is a more sophisticated version of the first one, meaning that individuals do not immediately copy any slightly superior solution (e.g. one that is only 1 percent better). Rather, agents only copy a solution if it is substantially better than one's own solution (e.g. 10 percent better). Following Norton *et al.* (2012), we label this as the Ikea-effect, based on the finding that even when individuals engage in a task as mundane as assembling Ikea boxes, they still prefer their own solutions compared to boxes of similar quality built by others.

We show that which behavioral assumption one implements in a simulation model substantially shapes the conclusions to be drawn about how to organize complex problem-solving. We replicate former studies showing that inefficient networks lead to better problem-solving performance (Barkoczi and Galesic, 2016; Lazer and Friedman, 2007) when relying on less empirically grounded assumptions where agents always copy the best solutions (either requiring a majority or not). However, when we implement the more realistic assumption that agents prioritize their own solutions, efficient networks perform better than inefficient networks. These results are consistent both when looking at average and maximum performance, although the Ikea-effect needs to be stronger when optimizing the maximum outcome.

Our study is thus in line with the micro-foundations approach increasingly advocated in strategy and management research, where insights into individual micro-foundations are argued to be necessary to understand macro-level collective behavior (Felin, Foss, and Ployhart, 2015; Puranam et al., 2015; Smith and Rand, 2018). We contribute to and integrate two influential streams of research: simulation studies on how to organize problem-solving (macro-level, e.g. March, 1991; Lazer and Friedman, 2007; Vuculescu and Bergenholtz, 2014; Schilling and Fang, 2014) and empirical studies of individual problem-solving behavior (micro-level, e.g. Billinger, Stieglitz, and Schumacher, 2014; Laureiro-Martinez et al., 2015; Levine, Reypens, and Riedl, 2017; Vuculescu, 2017). Rather than solely focusing on the network structure of individuals and assuming rationalized behavior, we directly compare the effectiveness of different individual search strategies, when individuals are embedded in different network structures. While we, in line with simulation studies, acknowledge that collective problem-solving requires some form of inefficiency to allow exploration of the solution space, we rely on empirical micro-level studies to show that this inefficiency does not necessarily need to be manifested at the structural level through inefficient network structures (Lazer and Friedman, 2007). An individual search bias that limits the imitation of all superior solutions implies friction at the individual level, just as inefficient networks create structural friction to ensure diversity. Metaphorically speaking, inefficiency is embodied at the individual level.

From a practitioner's perspective, our study has relevant implications for how to organize collective problem-solving, where balancing individual and social learning constitutes an important issue. Managers can influence to what degree problem solvers can learn from others or explore own solutions by varying the access to other problem solvers or by means of incentive systems that influence solvers' search behavior. This, for instance, has direct implications for the design of crowdsourcing platforms or innovation contests. While the majority of platforms do not allow solvers to see others' solutions (Wooten and Ulrich, 2015), our results illustrate that individuals *can* be exposed to the solutions of others, since an Ikea-effect characterizes the default search behavior of individuals, Moreover, when developing incentive systems, managers should think about not incentivizing individual solvers to immediately imitate others, since this would discourage explorative behavior. This is particularly important for problem-solving setups where one is searching for a maximum outcome and not just for the average.

In the next section we present the theoretical background, drawing upon both empirical and simulated studies of collective problem-solving in order to identify why and how inefficiency can be useful for different kinds of outcomes. Here after we document the setup of our simulation approach and present and analyze the results of our simulations. Finally, we discuss the implications and limitations of these findings, present managerial implications and outline an avenue for future research.

Collective Problem-Solving

Since the publication of March's seminal paper (1991), growing interest has been placed on the trade-off between exploration and exploitation. The risk with overemphasizing exploitation of known solutions is that problem solvers may get stuck into suboptimal solutions, thus not allowing sufficient time for exploration of more novel ones. A range of studies have focused on the possibility of using network structure as a lever to balance the exploration of novel solutions and the exploitation of promising known patches in the landscape, by varying the access to other problem solvers (Barkoczi and Galesic, 2016; Derex and Boyd, 2016; Fang et al., 2010; Lazer and Friedman, 2007; Mason and Watts, 2012; Schilling and Fang, 2014). The question is then how efficient, in terms of information flow, networks ideally should be.

If everyone has access to the solutions of everyone else in a given collective, promising solutions can be distributed quickly. Such a network is categorized as being efficient, since information flow is not limited. In contrast, if individuals only have access to a fraction of other problem-solvers, they spend more time on individual search for novel solutions, thus facilitating more exploration. Such a network is categorized as inefficient, since information flow is limited (Lazer and Friedman, 2007). While copying the best solutions facilitates exploitation, it also draws diversity out of the system, since the same current-best solutions are copied by multiple individuals. Yet, there is still disagreement on what types of social networks will benefit collective problem-solving performance the most and what underlying mechanisms drive the results (Bernstein, Shore, and Lazer, 2018; Derex and Boyd, 2016; Girotra, Terwiesch, and Ulrich, 2010; Lazer and Friedman, 2007; Mason and Watts, 2012; Schilling and Fang, 2014).

An important feature of these simulation studies is that they tend to operate under the implicit assumption that individuals always rely on superior social information when available. Some studies allow agents to copy any superior solution available in the collective (Kavadias and Sommer, 2009; Laland, Odling-Smee, and Myles, 2010; Lazer and Friedman, 2007; Posen, Lee, and Yi, 2013; Schilling and Fang, 2014), while others specify that agents imitate only when there is a majority occupying the superior solution (Barkoczi and Galesic, 2016; Fang *et al*., 2010; Rodan, 2008). The implication of this rationalized assumption is that friction can only be found at the structural level: Since superior solutions are always copied by agents in the system, one has to implement inefficiency at the network level in order to slow down a premature sharing of good solutions. So far, studies that have moved beyond pure imitation of superior solutions have relied on adding (random) errors to the copying process (Derex, Perreault, and Boyd, 2018; Kavadias and Sommer, 2009; Knudsen, Levinthal & Winter 2014), in order to introduce friction at the individual level. Yet, as noted by Puranam et al. (2015), simulation studies should be informed by empirical observations of actual search processes.

Even though copying the best and the majority are prevalent search strategies for largely conformist individuals, empirical evidence shows that individuals do not indiscriminately copy any superior solution[1] (Derex et al., 2015; Heck et al., 2018). Rather, individuals tend to follow a wide range of sophisticated social learning strategies with respect to when, what and whom to copy (Kendal et al., 2018), which are not based on a random choice as assumed by studies such as Kavadias and Sommer (2009). Furthermore, both lab and field studies in a variety of disciplines such as management (Hargadon and Bechky, 2006), psychology (Eriksson and Strimling, 2009), economics (Norton, Mochon, and Ariely, 2012; Weizsäcker, 2010), cultural evolution (Acerbi, Tennie, and Mesouidi, 2016; Derex et al., 2015; Mesoudi, 2011; Wisdom, Song, and Goldstone, 2013) as well as insights into the behavior of citizen scientists (Heck et al., 2018) clearly document that individuals have a tendency to prioritize their own solutions over social information.

Many different constructs are related to this behavior – e.g. over-confidence (Kahneman and Tversky, 1982), status-quo bias (Baumann and Martignoni, 2011), self-efficacy (Bandura, 1977), ego-centric discounting (Yaniv and Kleinberger, 2000) and feeling of ownership (Marsh et al., 2018). The discussion by Baumann and Martignoni (2011) is arguably the closest to our study, albeit at an organizational level. The authors argue that there are a number of forces that bias an organization towards over-evaluating its current solutions e.g. most notably organizational myopia (March 1991), but interestingly they argue that there are also opposing forces which put a premium on innovation: thus, a pro-innovation bias (Baumann and Martignoni 2001). Since our study aims to contribute to the discussion on micro-foundations of search, we prefer the construct 'Ikea-effect' (Norton et al. 2012) which speaks to an individual level cognitive bias. The argument is that individuals tend to value their own solutions more than solutions constructed by others as a consequence of the effort invested in developing them (Franke et al. 2010; Norton et al. 2012).

[1]Even in the original Asch experiments (Asch, 1955), only about 35 percent conformed in most conditions, and a stream of recent studies point to the limits of conformity (Eriksson and Coultas, 2009; Heyes, 2016; Mercier and Miton, 2019).

Norton et al. (2012) categorize this as the Ikea-effect, since they showed this effect in an unexciting context as building Ikea-boxes; even when the boxes were qualitatively identical, individuals preferred their own constructions compared to boxes of others. This individual bias is based on the endowment effect investigated by prior research (Pierce et al., 2001), which shows that individuals tend to attribute greater value to objects that they own. As such, this effect implies that individuals do not always and immediately copy any superior or most frequent solution, in contrast to assumptions in the literature. Rather, individuals prioritize their own solutions and are more likely to copy only if their own performance is exceeded substantially (Morgan et al. 2012; Derex et al. 2015).

Following Puranam et al. (2015) as well as Smith and Rand (2018) we argue that this calls for a stronger integration between empirical studies and assumptions used in simulations. We thus question the assumption that individuals engage in social learning indiscriminately (see e.g. Lazer and Friedman 2007) when they are observing peers' better performance and propose an agent-based approach where this individual bias (labelled the 'Ikea-effect') is taken into account. In our model this translates into individuals penalizing social information, which has to showcase a substantially better performance, in order to be copied by the individual. By doing so, we aim to show that the kinds of search behaviors solvers engage in can have a different influence on problem-solving performance. This has important implications for how to organize collective problem-solving.

In terms of the outcome, previous simulation work has largely focused on modeling systems where the average performance is the variable of interest (see e.g. Gavetti and Levinthal, 2000; Lazer and Friedman, 2007). However, as Boudreau et al. (2011) show, in actual innovation contests there are different optimal governance configurations for average vs. maximum performance. Innovation managers and contest organizers may, in fact, especially be interested in identifying one highly novel solution – e.g. an extreme value-outcome (Boudreau et al., 2011), even if this implies that the average quality of contributions decreases (Wang et al. 2018). This suggests that previous simulation results should be revisited in the context of innovation contests in particular and innovation management in general.

Methods

Fitness landscapes

We develop a NK model (Kauffman, 1993) that simulates the behavior of agents engaging in individual and social learning while solving a complex problem. The model is widely used in management research (Baumann, Schmidt, and Stieglitz, 2019) due to its simplicity in capturing the metaphor of search across a rugged landscape, with multiple suboptimal solutions distributed across it (Levinthal, 1997). A key challenge for problem-solvers is not to get stuck in suboptimal solutions when searching for the optimal performance. We study both simple environments with a global optimum (K=0) and complex environments with several suboptimal solutions (K=6,19). Formally, NK models can be thought of as optimization problems where the solution to the problem is represented as a vector of binary variables. In our case, we have set N=20. Then, the fitness function f is given by:

$$f(x)=\frac{1}{N}\sum_{i=1}^{N} x_i \qquad (1)$$

where x_i are the contributions of each of the 20 binary variables that form the vector. Contributions are drawn from a uniform distribution between 0 and 1, as follows: for K=0, every variable (or loci) has two possible contributions, corresponding to the two values the variable can take: 0 or 1. The contributions are randomly drawn from the distribution and remain fixed throughout the simulation. For K>0, the model posits there are K other variables influencing the contribution of any variable (Kauffman, 1993). Thus, for every 2^{k+1} possible combinations of 1s and 0s, a different contribution will be drawn from the uniform distribution. The landscapes used in this analysis are generated under the assumption that for any variable

x_i the influencing variables are its K consecutive neighbors, following the assumption that the vector is in fact a torus, so for K=6 the contribution of the variable x_{20} is influenced by $x_1, x_2 \ldots x_6$. To account for the stochasticity of the process, we generated 1000 different NK landscapes for K=0 and K=6 each. For K=19, only 400 such landscapes were generated, due to the high computational cost. We mapped the entire landscape rather than generating the landscapes dynamically and only computing fitness values as they are encountered by agents. Although faster, this latter approach only allows inference on average performance (Valente, 2008). In contrast, mapping the entire function allows us to identify which agents reached the maximum. As a further robustness check, following previous literature (Lazer and Friedman, 2007), we checked that the amplitude of the peaks in the landscape does not influence our results.

Agents and baseline model

In each simulation, 100 agents start from distinct positions in the landscape and attempt to find the maximum point. Agents therefore attempt to find the solution for the same combinatorial problem. They are connected in a network, whereby they have access to the solutions previously found by the other agents they have a direct connection with. Consistent with previous simulation models (Barkoczi and Galesic, 2016; Lazer and Friedman, 2007), we assume that an agent first engages in social learning behavior in order to find a better solution within its network and, only if unsuccessful, in individual learning.

Stopping rule, rounds

The agents are updated sequentially in each round. There is a maximum of 100 rounds in each simulation and for each experiment, the simulations are run 1000 times, once for each of the landscapes generated. All results are thus averaged over 1000 runs. A simulation is halted if the number of unique solutions in the population of 100 agents is 1 (Lazer and Friedman, 2007).

Copy the best strategy

For the baseline model, we model social learning as perfect imitation of the best solution found in the network, provided the given solution is superior compared to the previous solution (Lazer and Friedman, 2007). Alternatively, the agent attempts local search in the vicinity of the current solution (Lazer and Friedman, 2007), modeled as randomly selecting one of the 20 variables, changing its value from 0 to 1 or 1 to 0 and recomputing the fitness of the newly obtained solution. The agent 'moves' to the new location if and only if the fitness is superior to its previous one.

Ikea-effect strategy

A premium is added to a focal agent's performance before evaluating other solutions. The premium is a percentage of the maximum possible performance in the given landscape, since the best possible performance is landscape dependent due the stochastic nature of the fitness function. The baseline is set at 10 percent. However, we also test other percentages to analyze how different levels of the bias will influence search dynamics (see Figure 4).

Conformity strategy

Following Barkoczi and Galesic (2016) we here require that at least two agents have the given configuration of variables and that in situations where two different solutions are equally frequent, no social learning occurs. We also follow Barkoczi and Galesic (2016) and only report results for networks where agents are connected with three or more agents. We thus rely on the cited authors' operationalization of efficient vs. inefficient networks, which are modeled as random networks or fully connected networks where agents in a given round have access to a max number of *S* peers, the samples reporting being S=3 and S=9. However, given our high computational power we follow the original simulation study (Lazer and Friedman, 2007) and rely on S=100 as our fully connected network.

Networks

In line with previous studies (Barkoczi and Galesic, 2016; Fang *et al.*, 2010; Mason and Watts, 2012; Lazer and Friedman, 2007), we vary which network structures agents are embedded in, as these influence the speed of information diffusion. For copy the best and the Ikea-effect, we rely on three different types of networks: fully connected (agents are connected with 100 peers), ring lattice (agents are connected to 2 peers) and random graph (an average of 12 peers). As explained in the section above, for conformity simulations we only rely on random and fully connected networks. In the fully connected scenario, all agents can access and assess all other solutions within the same time step. In the lattice scenario, agents are connected with two other peers throughout the simulation. Finally, the random graph is generated using the algorithm in Watts and Strogatz (1998). The algorithm begins by constructing a regular ring lattice and with probability β every edge is rewired randomly such that loops and link duplications are not allowed. For β=1, this generates a random graph. Agents in the random graph have on average 12 peers (Watts and Strogatz, 1998). The network was kept constant throughout the simulation. We have thus chosen to focus on archetypical network structures. The lattice and the fully connected network structures constitute two extremes in terms of average path lengths, while random networks are still classified as inefficient, albeit with shorter average path lengths than lattices.

Data and code availability

Simulations were run in Matlab 2015b using the Abacus 2.0 supercomputer. The simulation code is available here: https://github.com/vulk29/ikea_search. Due to the large size of the NK landscapes (1 GB), these are not available for open download from Github but are available on request.

Results

The Ikea-effect and average performance in efficient and inefficient networks

Following previous literature (Barkoczi and Galesic, 2016; Lazer and Friedman, 2007), we first define problem-solving performance as the population's average standardized performance. In Figure 1 we compare the average performance achieved over time in both an efficient (fully connected) and two inefficient (either linear or randomly connected) networks for a complex environment (N=20, K=6), relying on the three different social learning strategies. As expected, we replicate former results (Lazer and Friedman, 2007; Barkoczi and Galesic 2016): on average, inefficient networks significantly outperform efficient networks when relying on the copy the best strategy (Figure 1a) or the conformity strategy (Figure 1b). However, when we implement an Ikea-effect (Norton et al., 2012), where agents have a bias toward their own solutions and only copy if other solutions are 10 per cent better than their current best, the search dynamics are completely different. Figure 1b shows that Ikea agents in efficient (fully connected) networks gain an increase in performance that flips previous results (Lazer and Friedman, 2007; Barkoczi and Galesic 2016), meaning that efficient networks outperform inefficient ones. This illustrates our main result: while non-Ikea-agents are dependent on friction at the structural level to maintain diversity, the Ikea-effect introduces friction at the agency level[2].

---- Insert Figure 1 here -----

The Ikea-effect and maximum performance in efficient and inefficient networks

Although previous literature has relied on average performance when studying collective problem-solving (Lazer and Friedman, 2007; Mason and Watts, 2012; Schilling and Fang, 2014), the use of this performance indicator may be misleading when dealing with new, unique problems. Average performance is a meaningful indicator when measuring performance related to standard tasks and routines. In contrast, when generating ideas or solving complex problems the aim is to search for the best possible solution and it is inconsequential whether one, several or all agents find the best solution. In line with studies on contests and innovation (Boudreau *et al.,* 2011; Girotra *et al*., 2010), we therefore change the focus to the maximum performance of *any* agent within the population. In Figure 2 we illustrate the evolution of maximum performance achieved over time by Ikea and non-Ikea agents (conformity and copy the best strategies) in

[2] Adding the Ikea-effect in both types of networks for a simple environment (K=0) and for a complex environment (K=19) leads to qualitatively similar results: introducing an Ikea-effect changes the search dynamics and leads to efficient networks outperforming inefficient ones. Note that in the case of simple problems, the conclusion is based on the speed with which the global maximum is found since all networks eventually find the maximum.

both efficient (fully connected) and inefficient (linear or randomly connected) networks for a complex environment (N=20, K=6). Results show how the search dynamics change significantly when implementing the Ikea-effect, where agents do not immediately copy any superior solution. The search dynamics are thus comparable to simulations where the average effect is the dependent variable. However, the strength of the Ikea-effect has to be substantially stronger when optimizing for the maximum outcome, compared to the average. We illustrate this by considering two levels of the Ikea-effect: 10 per cent and 40 per cent, chosen as proof-of concept. Note that Figure 4 illustrates that e.g. 30 or 50 percent yield similar overall performance as 40 percent. Given low levels of Ikea and copy the best strategies, inefficient networks outperform efficient networks. For higher levels of Ikea this result is overturned, and the performance is comparable or (marginally) higher for efficient networks. The same result is achieved by resorting to conformity social learning. This is because requiring that at least two agents have the same superior solution, in a large sample (S=100), is a relatively strong constraint and, as evidenced by Figure 3, even when fully connected, agents engage predominantly in local search.

---- Insert Figure 2 here -----

Mechanisms behind the Ikea-effect

After having compared the average and maximum performance achieved by Ikea and non-Ikea agents (i.e. agents following the copy the best or copy the majority strategy), in this section we focus on identifying the mechanisms that shape the Ikea-effect. How does the Ikea-effect influence the simulation dynamics? Previous studies point to the importance of maintaining some diversity in the system, since this facilitates an exploration of the search space that enhances collective problem-solving performance (Barkoczi and Galesic, 2016; Lazer and Friedman, 2007). As such, the number of unique solutions constitutes a key indicator for problem-solving performance. In Figure 3, we illustrate how adding the Ikea-effect impacts the diversity of solutions in the population of solvers, for a complex environment (N=20, K=6). These illustrations show that the Ikea-effect reduces the likelihood that agents converge prematurely on less optimal solutions. At the beginning of the simulation, efficient (fully connected) Ikea agents reach fewer unique solutions in the population compared to copy the best in inefficient networks. However, the average diversity for Ikea agents is superior to copy the best agents (see Figure 3a). Turning towards conformity strategies, it seems that conformity agents engage in far too little social learning (Figure 3b) as the number of unique solutions decreases very slowly over time. We thus can explain the relative inefficiency of conformity strategies through a diversity trap: as illustrated by Figure 3a nothing but the very worst solutions are eliminated from the population, which affects the average performance. This is because in the early stages of the simulation, the chances of finding at least two identical better solutions are quite small. This is also the explanation behind the counter intuitive fact that when considering maximum performance, conformity strategies perform better than copy the best strategies (see Table 1).

---- Insert Figure 3 and Table 1 here -----

Figure 3b illustrates the striking difference in evolution between the extent of local search in purely imitating agents vs. Ikea agents. In a population of non-Ikea agents (i.e. following the copy the best or copy the majority assumption), the exploration of the search space happens markedly different. Compared to the copy the best strategy, there is significantly more explorative search of the solution space since agents do not converge early on the current best solution, which may be a local peak[3]. Compared to the conformity strategy, Ikea agents engage in more social learning early on, which facilitates a more focused exploration in the later periods. In other words, allowing agents to be biased enables a better balance between social learning and exploration and enables a sort of simulated annealing mechanism (Van Laarhoven and Aarts,

[3] The two graphs look qualitatively the same when comparing inefficient networks.

1987) that gradually reduces solution diversity over time[4]. See Table 1 for further details on the ratio of social vs. individual learning in various configurations of search and networks.

Limits of the Ikea-effect

In this section we focus on comparing different levels of the Ikea-effect, not across different search strategies. Not surprisingly, there is a limit to the performance gains obtained by increasing the Ikea-effect. Figure 4 illustrates that the limit depends on the measure of performance chosen. While for average performance the optimal Ikea level is moderately low (for K=6, 10 per cent), the peak is reached at considerably higher Ikea levels for maximum performance (for K=6, 40 per cent). The difference between the two curves is due to 'stubbornness' being advantageous with respect to the exploration of the search space needed to find the best possible solution, while to improve average performance it is important to allow poorly performing agents to imitate higher performers, even at the risk of premature convergence. Also, it is important to highlight that, while these results arise as a consequence of facilitating more local search, local search alone is not a recipe for success, even in the case of 'smooth' problems. The reason is that local search cannot escape local traps in the case of moderately complex problems such as K=6 (Figure 4), and it is too slow in the case of a simple problem such as K=0.

---- Insert Figure 4 here -----

Discussion

When organizing for collective search the key challenge is to balance the exploration of novel solutions and exploitation of current best insights (March, 1991). Previous studies have been based on the argument that network structures should be inefficient, since this allows time for adequate exploration of the search space, reducing the risk of getting prematurely stuck in suboptimal solutions due to narrow exploitation

[4] The figure also illustrates the particular nature of simulations that assume imitating agents stop searching when there is only one solution left in the system; for the copy the best strategy, in an efficient network the simulation stops almost immediately.

(Bernstein *et al.,* 2018; Derex and Boyd, 2016; Girotra *et al.,* 2010; Lazer and Friedman, 2007; Schilling and Fang, 2014). We do not dispute that inefficiency is necessary to maintain solution diversity, yet compared to former simulations we argue that the tension is not merely a structural one and that assumptions about individual level behaviors are influential. Our point of departure is that simulation studies have assumed that individuals always switch to an available superior solution: Either agents immediately copy the best solution (Lazer and Friedman, 2007; Schilling and Fang 2014) or require a majority as a further necessary condition (Barkoczi and Galesic, 2016; Fang *et al.,* 2010). The implication of this rationalized assumption (Smith and Rand, 2018) is that friction needs to be introduced at the structural level to facilitate exploration.

Following Smith and Rand (2018) and Puranam *et al.* (2015) we argue that simulation studies should be directly informed by empirical studies of actual search behaviors. Although research does support that individuals have a conformist inclination towards superior solutions (Morgan et al., 2012), an in-depth reading of empirical studies on individual and social learning across disciplines reveals that individuals actually prioritize their own information over superior information available from others (Acerbi *et al.,* 2016; Derex et al., 2015; Eriksson and Strimling, 2009; Morgan et al., 2012; Norton *et al.*, 2012; Weizsäcker, 2010; Wisdom *et al.,* 2013). In fact, individuals only copy superior solutions if they are substantially better, which we label the Ikea-effect (Norton *et al.,* 2012). Based on these empirical insights, we carry out a systematic comparison of different individual search strategies, embedded in different network structures. Our main finding is that, when implementing the more realistic Ikea-effect, the search dynamics are fundamentally changed: Efficient networks now outperform inefficient networks. The theoretical mechanism driving this is that allowing individual agents to explore solutions that are inferior to the current best maintains diversity in the system, leading to overall improved performance. This shows that inefficiency does not have to be implemented at the structural level since an inherent human bias, the

Ikea-effect, embodies inefficiency at the individual level. More specifically, the friction that inefficient networks imply, can be manifested via an individual search bias that does not immediately copy any superior solution. Therefore, what could be categorized as an inefficient, individual level bias leads to macro-level efficiency and to a simulation dynamic where individual's propensity to search and explore rather than copying is the main driving force of optimal performance.

More generally, our model illustrates how one can follow Smith and Rand's (2018) proposal to integrate empirically informed micro-foundations of the collective system, without abandoning the simplicity a simulation model inherently requires (Harrison et al., 2007). Connecting and integrating macro-level simulation studies with micro-level empirical studies of individual problem-solving behaviors can thus open up new avenues for management research on how to organize collective problem-solving. Rather than focusing on implementing inefficiency at the network level, future research could focus on different organizational search contexts, incentives or other individual level factors as well as how micro- and macro-level features might interact.

We further develop the main contribution by acknowledging and contrasting how different sub-fields within management have different conceptions of what the appropriate outcome measure is. While literature on routines have focused on the collective's average performance as the primary outcome (Nelson, 2009), an innovation manager is likely to emphasize the best possible solution rather than the average (Boudreau *et al.,* 2011; Girotra *et al.,* 2010; Terwiesch and Xu, 2008). To our knowledge, we are the first to implement this distinction in a simulation of search performance in a collective. Our results turn out to be overall consistent across the two different outcome measures, however the Ikea-effect should be substantially stronger when one optimizes search for the best possible solution in the landscape (cf. Figure 4). When organizing for innovation, networks should thus either be more inefficient compared to organizing for routine tasks, or the organization should ensure less imitative behaviors of the individuals involved in the collective.

While the above contributions outline the overall Ikea-effect across network structures, it might be objected that non-Ikea agents could still be superior within network structures prevalent in the real world. However, results show that Ikea agents tend to outperform or are on par with non-Ikea agents in all cases apart from the lattice structure. Studies of real-world networks typically reveal small-world characteristics which substantially reduce friction at the structural level (Newman, 2001; Watts, 1999, 2004), illustrating that lattice structures are at best rare. Thus, while empirical studies highlight that individuals in fact have a bias towards their own solutions, our simulations illustrate that at a collective level, it is *beneficial* that individuals favor their own solutions, except for the most extremely inefficient situations.

Conclusion

Theoretical Contribution

As outlined above, our study has theoretical implications on how to balance inefficiencies at the individual and structural level, in order to enhance the organization of collective problem-solving. More broadly, our study thus contributes to and extends the problem-solving perspective of the firm (Nickerson and Zenger, 2004; Felin and Zenger, 2014). One of the claims of this perspective is that types of problems (with respect to e.g. complexity or hidden knowledge) should be matched with the mode of organization. Implicit to this perspective is the assumption that the search behaviors solvers engage in, is determined by the governance mechanisms set up by the focal organization. Incentives are in this context assumed to be embedded within specific organizational structures – organizations are likely to have low-powered incentives, which promote knowledge sharing, while e.g. platforms are assumed to have a mixed incentive system where intrinsic and extrinsic motives are essential (Felin and Zenger, 2014). Our results complement this perspective and suggest that even a low powered incentive mode is not enough. Rather one needs to take into account the kinds of search behaviors solvers are *prone* to engage in before deciding on the mode of organization.

Practical Implications

Our findings can provide relevant insights to managers with respect to the design of innovation contests. A key design aspect relates to moderating entries to such contests. Most platforms do not allow solvers to see each other's solutions to avoid a premature convergence toward sub-optimal solutions. However, our findings indicate that this is less of a concern, considering the existence of an Ikea-effect with respect to individuals' search behaviors. Our findings thus highlight that individuals can be exposed to others' solutions. This is important because in such platforms, one often observes a 'rich get richer' phenomenon where platforms with a high level of visible activity generate more traffic, and platforms with less visible activity discourage further traffic (Wooten an Ulrich, 2015). Moreover, especially when aiming for maximum performance, managers should consider developing incentive systems that increase the Ikea-effect, rather than incentivizing for immediate coping of superior solutions from others. Managers should thus not discourage explorative behaviors, in order to allow individuals to push their own solutions in the system and maintain diversity.

Furthermore, we extend former managerial recommendations focusing on network-induced inefficiency (Lazer and Friedman, 2007; Schilling and Fang 2014). We replicate and acknowledge former simulation studies showing that network structures do matter. Temporarily constraining opportunities for social learning is likely to enhance search performance (Bernstein et al., 2018; Derex and Boyd, 2016; Lazer and Friedman, 2007). However, one should note that the Ikea-effect is a stylized, not a fixed fact, and that incentives influence search behavior (Ederer and Manso, 2013; Lee and Meyer-Doyle, 2017). For example, pay systems contingent on current performance can reduce individual exploration (Lee and Meyer-Doyle, 2017). Managers should thus consider what reward structure (e.g. no rewards, paying average performance, paying maximum performance) is the most appropriate solution for the given problem (complexity) and network structure (efficient or not) at hand, in order to balance the overall inefficiency originating in both network and individual level features.

Future Research

In simulations such as these, one tests the influence of the Ikea-effect 'all things being equal', meaning that factors not included in the model (e.g. individual status, expertise or incentive structures) may influence the size of the effect of this individual bias. While the Ikea-effect is empirically well established as a stylized fact, the effect is mainly based on experiments in a lab-setting, where students have acted as participants and have been rewarded based on the average performance of their search activities (e.g. Acerbi, Tennie, and Mesouidi, 2016; Derex et al., 2015, while Heck et al.'s (2018) focus on max performance is an exception). It is therefore an open question what the actual and optimal size of the Ikea-effect is across various incentive structures, complexities (K) of the problem task, as well as across individual heterogeneities such as cognitive ability (Glowacki and Molleman, 2017; Muthukrishna, Morgan, and Henrich, 2016) or expertise (Jeppesen and Lakhani 2010). People with high cognitive abilities have been shown to be better at assessing when it is effective to copy others (Vostroknutov, Polonio, and Coricelli, 2018), and rewarding people based on their maximum, not their average, performance should increase explorative behavior (Ederer and Manso, 2013). Since these factors are likely to shape search and copying behavior, the optimal network structure is likely to vary. Furthermore, one could imagine that in an organizational setting an opposing force such as an innovation bias (Baumann and Martignoni, 2011) could counter-act the Ikea effect. Overall, the field would benefit from further empirical studies that could strengthen the link between macro- and micro-levels of analysis (cf. Puranam et al. 2015; Baumann et al. 2019).

Furthermore, in our study, we focus on network structures in our simulation, not on individual network positions. It would be relevant to manipulate the degree to which individual agents occupy bridging structural holes or a central network position, in order to investigate how such changes might influence the search behavior and performance of individuals as well as the overall collective. In addition, our model mirrors Lazer and Friedman (2007) who implemented the copy-first, local-search-second rule, yet it would be interesting to examine how reverting this search sequence would affect both individual and collective search and performance. Finally, following previous studies we developed our model based on the widely used NK function in which we depict a complex problem as a rugged landscape. Future research should investigate whether our results apply to other types of landscapes, that have different properties aside from ruggedness, such as modularity (Egidi and Marengo, 2004) or deceptiveness (Vuculescu et al. 2020).

References

Acerbi, A., Tennie C., & Mesoudi, A. (2016). Social learning solves the problem of narrow-peaked search landscapes: experimental evidence in humans. *Royal Society Open Science, 3*(9), 160215.

Adamczyk, S., Bullinger, A. C., & Möslein, K. M. (2012). Innovation contests: A review, classification and outlook. *Creativity and Innovation Management*, *21*(4), 335-360.

Asch, S.E. (1955). Opinions and social pressure. *Scientific American, 193*(5), 31–35.

Bandura, A. (1977). Self-efficacy: toward a unifying theory of behavioral change. *Psychological review, 84*(2), 191.

Barkoczi, D., & Galesic, M. (2016). Social learning strategies modify the effect of network structure on group performance. *Nature Communications, 7*, 13109.

Baumann, O. & Martignoni, D. (2011). Evaluating the new: the contingent value of a pro-innovation bias. *Schmalenbach Business Review*, 63(4), 393-415.

Baumann O, Schmidt J, & Stieglitz N. (2019). Effective search in rugged performance landscapes: A review and outlook. *Journal of Management, 45*(1), 285–318.

Bernstein, E., Shore, J., & Lazer, D. (2018). How intermittent breaks in interaction improve collective intelligence. *Proceedings of the National Academy of Sciences, 115*(35), 8734–8739.

Billinger, S., Stieglitz, N., & Schumacher, T.R. (2014). Search on rugged landscapes: An experimental study. *Organization Science, 25*(1), 93–108.

Boudreau, K.J., Lacetera, N., & Lakhani K.R. (2011). Incentives and problem uncertainty in innovation contests: An empirical analysis. *Management Science, 57*(5), 843–863.

Boudreau, K. J., & Lakhani, K. R. (2015). “Open” disclosure of innovations, incentives and follow-on reuse: Theory on processes of cumulative innovation and a field experiment in computational biology.

Research Policy, *44*(1), 4-19.
Boyd, R., Richerson, P. J., & Henrich, J. (2011). The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences*, *108*(Supplement 2), 10918-10925.
Cooper, S., Khatib, F., Treuille, A., Barbero, J., Lee, J., Beenen, M., Leaver-Fay, A., Baker, D., Popovic, Z., et al. (2010). Predicting protein structures with a multiplayer online game. *Nature, 466*(7307), 756.
Derex, M., & Boyd, R. (2016). Partial connectivity increases cultural accumulation within groups. *Proceedings of the National Academy of Sciences, 113*(11), 2982–2987.
Derex, M., Feron, R., Godelle, B., & Raymond, M. (2015). Social learning and the replication process: an experimental investigation. *Proceedings of the Royal Society B: Biological Sciences, 282*(1088), 20150719.
Derex, M., Perreault, C., & Boyd, R. (2018). Divide and conquer: intermediate levels of population fragmentation maximize cultural accumulation. *Philosophical Transactions of the Royal Society B: Biological Sciences, 373*(1743), 20170062.
Ederer, F., & Manso, G. (2013). Is pay for performance detrimental to innovation? *Management Science, 59*(7), 1496–1513.
Egidi, M., & Marengo, L. (2004): Cognition, institutions, near decomposability: Rethinking Herbert Simon's contribution. *In Models of A Man: Essays in Memory of Herbert A. Simon, Part III*. Edited by Augier, M., March, J.J. (eds.), Cambridge MA: MIT Press.
Eriksson, K., & Coultas, J. (2009). Are people really conformist-biased? An empirical test and a new mathematical model. *Journal of Evolutionary Psychology, 7*(1), 5–21.
Eriksson, K., Strimling, P. (2009). Biases for acquiring information individually rather than socially. *Journal of Evolutionary Psychology, 7*(4), 309–329.
Fang, C., Lee, J., Schilling, M.A. (2010). Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organization Science, 21*(3), 625–642.
Franke, N., Schreier, M., & Kaiser, U. (2010). The "I designed it myself" effect in mass customization. *Management science, 56*(1), 125-140.
Felin, T., & Zenger, T. R. (2014). Closed or open innovation? Problem solving and the governance choice. *Research Policy, 43*(5), 914-925.
Felin, T., Foss, N.J., Ployhart, R.E. (2015). The microfoundations movement in strategy and organization theory. *Academy of Management Annals, 9*(1), 575–632.
Gavetti, G., & Levinthal, D. (2000). Looking forward and looking backward: Cognitive and experiential search. *Administrative science quarterly*, *45*(1), 113-137.
Girotra, K., Terwiesch, C., & Ulrich, K.T. (2010). Idea generation and the quality of the best idea. *Management Science, 56*(4), 591–605.
Glowacki, L., & Molleman, L. (2017). Subsistence styles shape human social learning strategies. *Nature Human Behavior, 1*(5), 0098.
Hargadon, A.B., & Bechky, B.A. (2006). When collections of creatives become creative collectives: A field study of problem solving at work. *Organization Science, 17*(4), 484–500.
Harrison, J.R, Lin, Z., Carroll, G.R., & Carley, K.M. (2007). Simulation modeling in organizational and management research. *Academy of Management Review, 32*(4), 1229–1245.
Heck, R., Vuculescu, O., Sørensen, J.J., Zoller, J., Andreasen, M.G., Bason, M.G., Ejlertsen, P., Elıasson, O., Haikka, P., Laustsen, J.S., et al. (2018). Remote optimization of an ultra-cold atoms experiment by experts and citizen scientists. *Proceedings of the National Academy of Sciences*, 115(48), E11231-E11237.
Heyes, C. (2016). Who knows? Metacognitive social learning strategies. *Trends in Cognitive Sciences, 20*(3), 204 213.
Kahneman, D., & Tversky, A. (1982). The psychology of preferences. *Scientific American, 246*(1), 160-173.
Kauffman, S.A. (1993). *The origins of order: Self-organization and selection in evolution.* OUP USA
Kavadias, S., & Sommer, S.C. (2009). The effects of problem structure and team diversity on brainstorming effectiveness. *Management Science, 55*(12), 1899–1913.
Kendal, R.L., Boogert, N.J., Rendell, L., Laland, K.N., Webster, M., & Jones, P.L. (2018). Social learning strategies: bridge-building between fields. *Trends in Cognitive Sciences, 22*(7), 651-665.

Jeppesen, L. B., & Lakhani, K. R. (2010). Marginality and problem-solving effectiveness in broadcast search. *Organization Science, 21*(5), 1016-1033.
Laland, K.N., Odling-Smee, J., & Myles, S. (2010). How culture shaped the human genome: bringing genetics and the human sciences together. *Nature Review Genetics, 11*(2), 137.
Laureiro-Martinez, D., Brusoni, S., Canessa, N., & Zollo, M. (2015). Understanding the exploration–exploitation dilemma: An fmri study of attention control and decision-making performance. *Strategic Management Journal, 36*(3), 319–338.
Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly, 52*(4), 667–694.
Lee, S., & Meyer-Doyle, P. (2017). How performance incentives shape individual exploration and exploitation: Evidence from microdata. *Organization Science, 28*(1), 19–38.
Levine, S.S., Reypens, C., & Riedl, C. (2017). Should I stay or should I go? The cognition of exploration and exploitation. *Working paper (November 11, 2017)*.
Levinthal, D.A. (1997). Adaptation on rugged landscapes. *Management Science, 43*(7), 934–950.
March, J.G. (1991). Exploration and exploitation in organizational learning. *Organization Science, 2*(1), 71–87.
Marsh, L. E., Kanngiesser, P., & Hood, B. (2018). When and how does labour lead to love? The ontogeny and mechanisms of the IKEA effect. *Cognition, 170*, 245-253.
Mason, W., Watts, D.J. (2012). Collaborative learning in networks. *Proceedings of the National Academy of Sciences; 109*(3), 764–769.
Mercier, H., & Miton, H. (2019). Utilizing simple cues to informational dependency. *Evolution and Human Behavior, 40*(3), 301-314.
Mesoudi, A. (2011). An experimental comparison of human social learning strategies: payoff-biased social learning is adaptive but underused. *Evolution and Human Behavior, 32*(5), 334–342.
Morgan, T., Rendell, L., Ehn, M., Hoppitt, W., & Laland, K. (2012). The evolutionary basis of human social learning. *Proceedings of the Royal Society B: Biological Sciences, 279*(1729), 653–662.
Muthukrishna, M., Morgan, T.J., & Henrich, J. (2016). The when and who of social learning and conformist transmission. *Evolution and Human Behavior, 37*(1), 10–20.
Nelson, R.R. (2009). *An evolutionary theory of economic change.* Harvard University Press.
Newman, M.E. (2001). The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences, 98*(2), 404–409.
Nickerson, J. A., & Zenger, T. R. (2004). A knowledge-based theory of the firm—The problem-solving perspective. *Organization Science*, *15*(6), 617-632.
Norton, M.I., Mochon, D., & Ariely, D. (2012). The Ikea effect: When labor leads to love. *Journal of Consumer Psychology, 3*(22), 453–460.
Nowak, S., & Krug, J. (2015). Analysis of adaptive walks on NK fitness landscapes with different interaction schemes. *Journal of Statistical Mechanics: Theory and Experiment, 2015*(6), P06014.
Pierce, J. L., Kostova, T., & Dirks, K. T. (2001). Toward a theory of psychological ownership in organizations. *Academy of Management Review, 26*(2), 298-310.
Posen, H.E., Lee, J., & Yi, S. (2013). The power of imperfect imitation. *Strategic Management Journal, 34*(2), 149–164.
Posen, H.E., Yi, S., & Lee, J. (2020). A contingency perspective on imitation strategies: When is “benchmarking” ineffective? *Strategic Management Journal, 41(2),* 198-221.
Puranam, P., Stieglitz, N., Osman, M., & Pillutla, M.M. (2015). Modelling bounded rationality in organizations: Progress and prospects. *Academy of Management Annals, 9*(1), 337–392.
Rodan, S. (2008). Organizational learning: effects of (network) structure and (individual) strategy. *Computational and Mathematical Organization Theory, 14*(3), 222–247.
Schilling, M.A., & Fang, C. (2014). When hubs forget, lie, and play favorites: Interpersonal network structure, information distortion, and organizational learning. *Strategic Management Journal, 35*(7), 974–994.
Smith, E.B., & Rand, W. (2018). Simulating macro-level effects from micro-level observations. *Management Science, 64*(11), 5405-5421.
Terwiesch, C., & Xu, Y. (2008). Innovation contests, open innovation, and multiagent problem solving. *Management*

Science, 54(9), 1529–1543.
Valente, M. (2008). Pseudo-nk: an enhanced model of complexity. Working paper.
Van Laarhoven, P.J., & Aarts, E.H. (1987). Simulated annealing. *Simulated annealing: Theory and applications*, 7–15. Springer.
Vostroknutov, A., Polonio, L., & Coricelli, G. (2018). The role of intelligence in social learning. *Scientific Reports, 8*(1), 6896.
Vuculescu, O., & Bergenholtz, C. (2014). How to solve problems with crowds: a computer-based simulation model. *Creativity and Innovation Management*, *23*(2), 121-136.
Vuculescu, O. (2017). Searching far away from the lamp-post: An agent-based model. *Strategic Organization, 15*(2), 242– 263.
Vuculescu, O., Petersen, MK., Sherson, JF., & Bergenholtz, C. (2020). Human search in a fitness landscape: How to assess the difficulty of a search problem. *Complexity* (published online).
Wang, K., Nickerson, J., & Sakamoto, Y. (2018). Crowdsourced idea generation: The effect of exposure to an original idea. *Creativity and Innovation Management*, *27*(2), 196-208.
Watts, D.J. (1999). Networks, dynamics, and the small-world phenomenon. *American Journal of Sociology, 105*(2), 493–527.
Watts, D.J. (2004). The new science of networks. *Annual Review of Sociology, 30,* 243–270.
Watts, D.J, & Strogatz, S.H. (1998). Collective dynamics of small-world networks. *Nature, 393*(6684), 440–442.
Weizsäcker, G. (2010). Do we follow others when we should? A simple test of rational expectations. *American Economic Review, 100*(5), 2340–2360.
Wisdom, T.N., Song, X., & Goldstone, R.L. (2013). Social learning strategies in networked groups. *Cognitive Science, 37*(8), 1383–1425.
Wooten, J. O., & Ulrich, K. T. (2015). The impact of visibility in innovation tournaments: Evidence from field experiments. Available at SSRN 2214952.
Yaniv, I., & Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational behavior and human decision processes, 83*(2), 260-281.

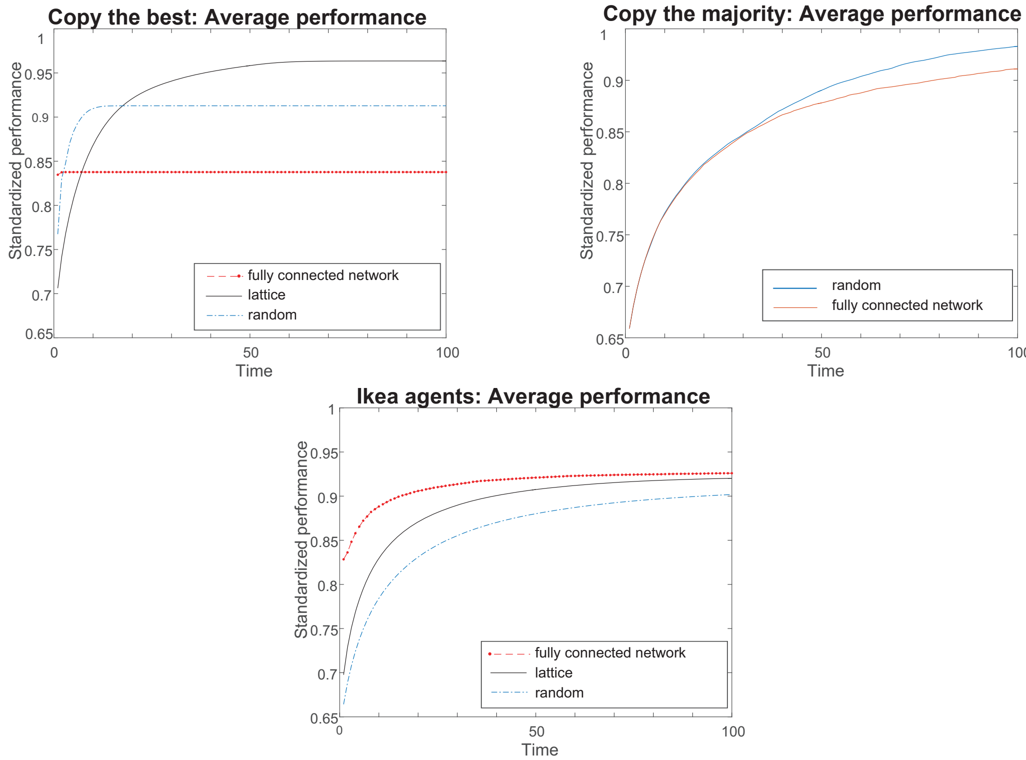

Figure 1. The Ikea-effect and average performance over time in efficient and inefficient networks. Results are reported for complex environments with multiple local optima and a global optimum (N=20, K=6) both for efficient and inefficient networks. The top left panel (a) illustrates the standardized average performance for agents who imitate the best solution. The top right panel (b) shows the standardized average performance achieved by agents who imitate the best solution as long as they encounter it at least twice in their network (conformity). The bottom panel (c) shows the standardized average performance achieved when an Ikea-effect (10 per cent level) is introduced. Performance is reported here as the ratio between average performance and the maximum possible performance in the given landscape. Blue dotted curves represent the performance of agents in randomly connected networks; Black solid curves represent in lattice networks (NB: one cannot implement a conformity strategy in a lattice network); Red dotted curves represent in fully connected networks.

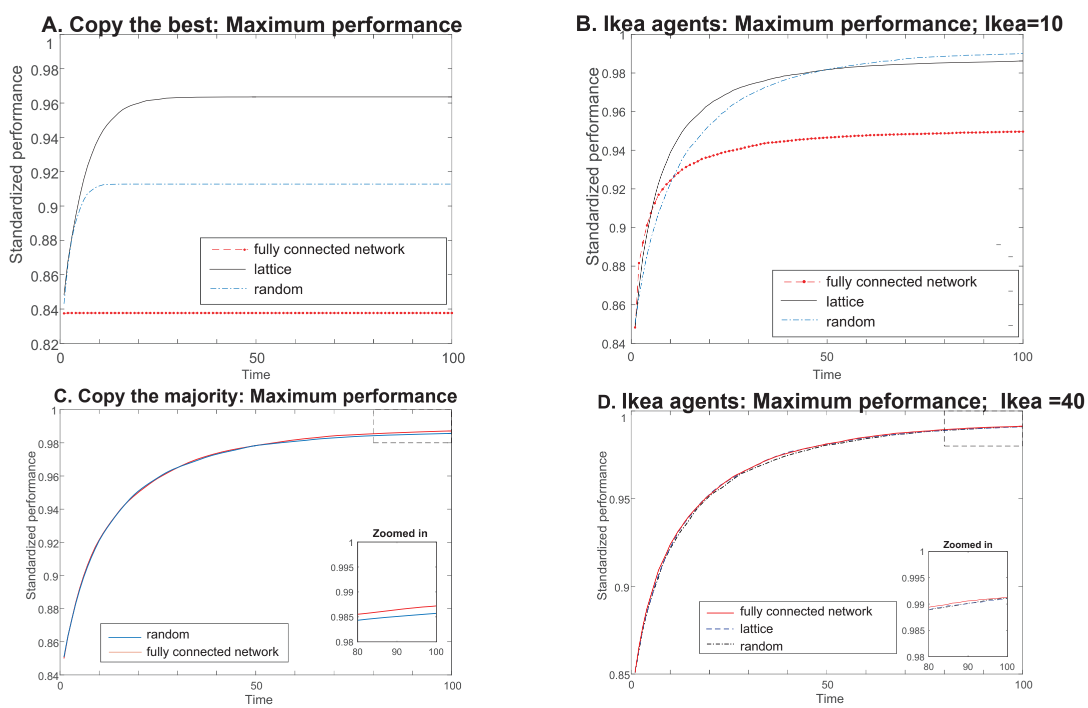

Figure 2. The Ikea-effect and maximum performance over time in efficient and inefficient networks. Results are reported for complex environments with multiple local optima and a global optimum (N=20, K=6) for efficient and inefficient networks. Maximum performance is computed as the ratio between the best performing agent at a given time and the maximum possible performance achievable in the given landscape. (a) maximum performance for non-Ikea agents; (b) maximum performance for Ikea agents when an Ikea-effect (10 per cent level) is introduced; (c) maximum performance for Ikea agents when an Ikea-effect (40 per cent level) is introduced; (d) comparison between non-Ikea and Ikea agents in terms of maximum performance achieved in both types of networks. Blue dotted curves represent non-Ikea (a) and Ikea agents (b, c) in randomly connected networks; Black continuous curves represent non-Ikea (a) and Ikea agents (b, c) in lattice networks; Red dotted curves represent non-Ikea (a) and Ikea agents (b, c) in fully connected networks; In panel (d), the red continuous curve represents Ikea agents in fully connected networks; Red dashed curve represents non-Ikea agents in fully connected networks; Blue dotted curve represents Ikea agents in lattice networks; and Black solid curve represents non-Ikea agents in lattice networks.

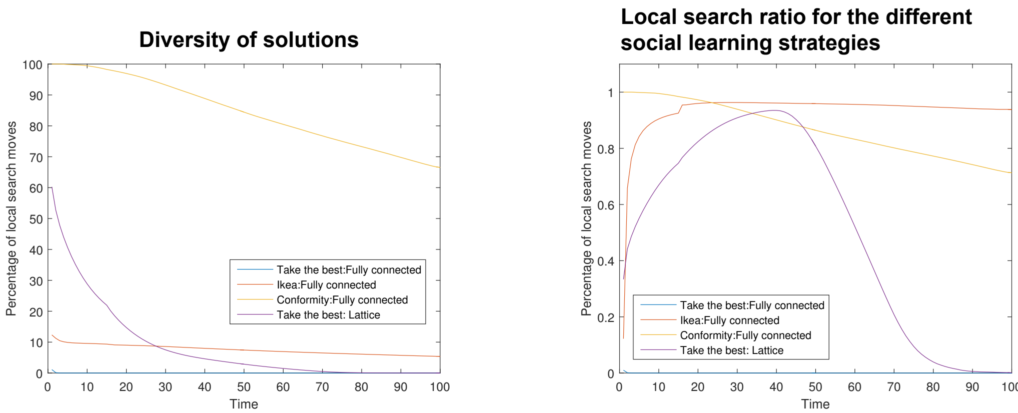

Figure 3.The Ikea-effect on solution diversity and local search. Results are reported for complex environments with multiple local optima and a global optimum (N=20, K=6) for all three social learning strategies with efficient networks. Take the best with inefficient networks was added as a benchmark. The left panel a) shows a direct comparison with respect to the number of unique solutions in a population of solvers at a given point. This number is averaged over 1000 runs. The right panel (b) shows the ratio of successful local search moves over time. Red line; the evolution of the local search moves ratio for Ikea 10. Yellow line; the evolution of the local search moves ratio for a conformity strategy within a fully connected network. Copy the best within a fully connected network, the blue line, hovers just above 0 and converges to 0 after time=3, since simulations stop when there is only one solution left. The purple line represents the evolution for a Copy the best strategy within a ring lattice network.

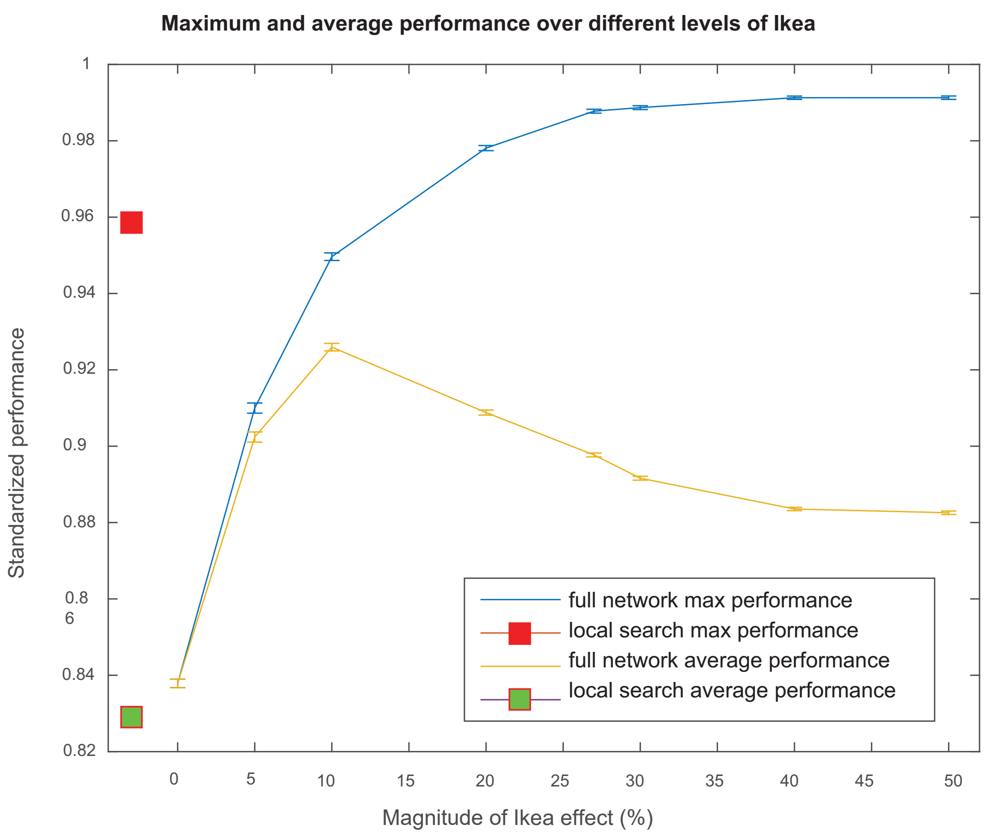

Figure 4.Limits of the Ikea-effect. Results are reported for complex environments with multiple local optima and a global optimum (N=20, K=6) for efficient networks. Maximum (blue line) and average (yellow line) performance for different magnitudes of Ikea percentages (x-axis). Local search is shown as a benchmark: the red square represents the standardized maximum performance achieved by an agent who only engages in local search (thus social learning penalties are meaningless in this context). The green square represents the standardized average performance achieved by an agent who only engages in local search (thus social learning penalties are meaningless in this context).

Social rule	Average (std)	Maximum	Local search ratio	Basins left	Speed	Diversity
Take the best (Full)	0.84 (0.04)	0.83	0.004	505	1.12	0.01
Take the Best (Lattice)	0.96 (0.03)	0.96	0.52	416	60.79	7.57
Take the best (Random)	0.91 (0.04)	0.91	0.14	489	7.53	0.45
Ikea (Full)	0.93 (0.03)	0.94	0.94	458	95.84	7.54
Ikea (Lattice)	0.92 (0.02)	0.98	0.98	258	99.96	53
Ikea (Random)	0.90 (0.01)	0.99	0.99	210	100	81.54
Conformity(Random)	0.93 (0.06)	0.98	0.87	153	88.67	70.75
Conformity(Full)	0.91 (0.05)	0.98	0.86	254	86.99	84.57

Table 1 Summary of simulation runs. Ikea effect reported here at 10 per cent. Average performance standardized and the numbers in the parenthesis (std) refer to the standard deviation. Local search ratio refers to how many of the searches are local vs. social, on a scale from 0 to 1. A local search is a successful local search move. Basins left refers to how often a basin that would lead to the global optimum was abandoned. Speed refers to how many rounds are needed in order to reach the maximum performance, averaged over 1000 runs. Average diversity refers to the number of unique solutions left in the system at the end of the simulation run.