Journal of Business Research 129 (2021) 589–599

Contents lists available at ScienceDirect

Journal of Business Research

journal homepage: www.elsevier.com/locate/jbusres

A network-based model of exploration and exploitation

Pieter den Hamer [a,b], Koen Frenken [c,*]

[a] Centre for Complex Systems Studies (CCSS), Utrecht University, the Netherlands
[b] Gartner, the Netherlands
[c] Innovation Studies, Copernicus Institute of Sustainable Development, Utrecht University, the Netherlands

ARTICLE INFO

Keywords:
Complexity
NK-model
Proximity
Small-world

ABSTRACT

We propose a new model of exploration and exploitation, in which firms rely on local search for exploitation and on imitation for exploration. We assume that firms imitate the knowledge base of successful competitors, with imitation errors taking place depending on the social distance between the imitating firm and imitated firm in the network. The key model outcome, consistent with earlier empirical findings, holds that successful imitation generally occurs at an intermediate level of cognitive proximity because imitation at high cognitive distance is too error-prone, while for imitation at low cognitive distance there are typically no firms to imitate. A second outcome holds that social and cognitive proximity are substitutes. The model further shows that exploration by imitation is more beneficial in highly complex industries than in less complex industries, and that small-world networks yield the highest benefits for collective learning.

1. Introduction

The distinction between exploration and exploitation has proven helpful in understanding how firms innovate, and the tensions they need to balance within and outside their boundaries (Lavie, Stettner, & Tushman, 2010; March, 1991). In the context of innovation, the distinction between exploration and exploitation is often mapped onto the notions of radical versus incremental innovation in firms, but in essence exploration and exploitation refer to two generic modes of learning at the level of individuals, teams, and organizations alike (Gupta, Smith, & Shalley, 2006; Wilden, Hohberger, Devinney, & Lavie, 2018). When reasoning about firms, exploitation can generally be well managed within the firm's boundaries as exploitation leverages existing knowledge as to incrementally improve a firm's activities and outputs (Benner & Tushman, 2003; March, 1991). By contrast, exploration involves a search for new and distant knowledge through recombination, experimentation and risk-taking (March, 1991; Savino, Messeni Petruzzelli, & Albino, 2017).

When both exploration and exploitation are undertaken in-house, conflicts tend to arise given that exploitation routines are so different from exploration routines (Stettner & Lavie, 2014). Instead, firms often engage in exploration by looking externally for new knowledge, for example, by imitating knowledge held by other firms (Csaszar & Siggelkow, 2010). Indeed, as empirical research has shown, imitation is a salient feature of firms' learning strategies (for a review, see Ordanini, Rubera, & De Fillippi, 2008).

It is common to assume that imitation is not a blind process, but biased towards successful competitors (Lieberman & Asaba, 2006; Nelson & Winter, 1982). What distinguishes imitation from other ways of engaging in exploration, then, holds that the *intended* outcome of exploration is well-defined (imitate the solution held by a successful competitor). However, imitation is difficult and failure-prone as the knowledge required to successfully imitate is generally quite distinct from the knowledge already present in a firm (Baumann, Schmidt, & Stieglitz, 2019). Especially in the context of high product complexity, small copying errors can lead to drastic reductions in performance (Rivkin, 2000). Given the original formulation by March (1991, p. 85) that exploration concerns "experimentation with new alternatives" with its returns being "uncertain" and "often negative", imitation can be considered a form of exploration. In an exploration–exploitation framework, then, one can view exploitation as involving local search, building on a firm's existing knowledge yielding predictable increments in performance, and exploration through imitation as a jump away from its existing knowledge, with uncertain and often negative results on performance (Csaszar & Siggelkow, 2010, p. 674).

The key question that follows holds what makes a firm successful in

* Corresponding author at: Innovation Studies, Copernicus Institute of Sustainable Development, Utrecht University, P.O. Box 80115, 3508TC Utrecht, the Netherlands.

E-mail address: k.frenken@uu.n (K. Frenken).

https://doi.org/10.1016/j.jbusres.2019.12.040
Received 8 February 2019; Received in revised form 22 December 2019; Accepted 23 December 2019
Available online 15 January 2020
0148-2963/

imitation. Or, more specifically, the question holds under what conditions firms can avoid making copying errors when they intend to imitate a superior solution of a competitor. In the past, the question of successful imitation has been primarily approached from two angles. One strand of literature looks at the firm's knowledge base by investigating what properties of a firm's knowledge base adds to its absorptive capacity (Cohen & Levinthal, 1990; Nooteboom, 2000). The more relevant knowledge a firm already possesses, the easier it will be to absorb new knowledge by imitation. A second strand of literature looks at the social relations between the imitating and imitated firm. Here, one can further distinguish between formal ties such as licenses that purposefully support imitation (Laursen, Leone, & Torrisi, 2010) versus informal social networks between employees of firms as channels for knowledge spillovers (Breschi & Lissoni, 2009).

The theory of exploration and exploitation we propose aims to combine the knowledge base perspective and the relational perspective in a single stylized model. We look at firms in different industry contexts characterized with different levels of complexity to understand whether the relative value of exploration and exploitation varies for products with different levels of complexity. While learning is especially difficult for firms dealing with complex products, innovation in such industries nevertheless relies a lot on inter-firm learning (Miller, Hobday, Leroux Demers, & Olleros, 1995; Powell, Koput, & Smith-Doerr, 1996). Firms that fully rely on their internal knowledge in innovating complex products are bound to end up in poor local optima (Levinthal, 1997). Hence, such firms would especially benefit from supplementing their internal innovation efforts (exploitation) with learning by imitating knowledge from others (exploration) as to escape such poor optima (Csaszar & Siggelkow, 2010). Given the importance of absorptive capacity, though, imitation is more likely to be successful if the knowledge bases of the imitating and imitated firms overlap considerably (Nooteboom, 2000).

Regarding social relations between the imitating and imitated firm, we are agnostic about the specific type of ties that would support imitation, nor do we want to limit our perspective to direct ties only. Instead, aiming for generalizability, we want to look at the effect of social networks on imitation attempts between any pair of firms, be them with direct ties (at social distance of 1) or indirect ties (at social distance larger than 1), while acknowledging that imitation will become more error-prone at longer social distances. In adopting this generalized social network perspective, we move beyond earlier exploration-exploitation models (Lazer & Friedman, 2007; Miller, Zhao, & Calantone, 2006) and related models on inter-firm learning (Cowan & Jonard, 2003, 2004, 2009; Cowan, Jonard, & Özman, 2004). In these models, firms only copy solutions through direct ties as it would occur in formal inter-firm strategic alliances, while in our model solutions can be imitated among any two firms which are all part of a single social network.

In our investigation, we will distinguish between two levels of analysis. First, at the network level, we will compare a range of network structures that differ in terms of the average social distance between firms and the average clustering between firms using the 'small-world' parameter (Watts & Strogatz, 1998). This parameter tunes the structure of a network from fully regular to 'small-world' to fully random. The small-world network combines the feature of high clustering of a regular network with that of short distances of a random network. Comparing these networks allows us to investigate whether high clustering and short distances are complementary for imitation, as found in other models where firms only learn direct partners (Cowan & Jonard, 2003, 2004) and in empirical studies on the role of networks on innovation (Capaldo, 2007; Fleming, Chen, & Mingo, 2007; Schilling & Phelps, 2007). Second, at the level of dyads concerning each pair of firms, we investigate whether successful instances of imitation between two firms occur at an intermediate cognitive distance as implied by the thesis of optimal cognitive proximity (Cowan & Jonard, 2009; Nooteboom, 2000). We further analyze whether socially proximate firms may be better able to learn cognitively distant knowledge as compared to socially distant firms, that is, whether social proximity can compensate for cognitive distance (Boschma, 2005; Huber, 2012).

Our contributions, then, are two-fold. First, in the analysis of exploration by imitation, we integrate insights from absorptive capacity, social network and complexity theories into a single modelling framework. We analyze how effective imitation between firms is affected by both cognitive distance and social distance, while also taking into account the complexity of the knowledge base at hand. Second, we aim to reproduce a diverse set of empirical findings regarding (i) the existence of an optimal level of cognitive proximity in imitation, (ii) the substitution effect between cognitive and social proximity, (iii) the high benefits of exploration in complex-product industries compared to simple-product industries, and (iv) the benefits of small-worlds for collective learning,

2. Theory

In processes of learning, it is customary to distinguish between exploration and exploitation. In this view, firms learn both by exploiting their existing knowledge and by exploring new knowledge (March, 1991). While exploitation activities build closely on a firm's internal knowledge, exploration activities often rely on knowledge found outside a firm's own organization. In particular, firms have the tendency to imitate better performing competitors (Nelson & Winter, 1982). In this sense, the imitation of better performers underlies the evolutionary logic of markets ensuring the diffusion of superior solutions at the expense of inferior solutions in a population of competing firms (Lieberman & Asaba, 2006).

Imitation, however, should not be equated with a simple copying process among firms (Nelson & Winter, 1982). Imitation attempts are prone to errors, as firms may struggle to correctly interpret knowledge from others (Rivkin, 2000). The efforts involved in learning from other firms may be in vain if such attempts result in only partial understanding with limited economic return. In particular, the effectiveness of inter-firm learning depends on the complexity of knowledge to be learnt (Rivkin, 2000). The complexity of a product, a technology or service can be judged from the number of interdependencies between the components that make up a product, technology or service (Gatti, Volpe, & Vagnani, 2015; Levinthal, 1997; Simon, 1996 [1969]). High complexity requires finely tuned component assemblies to yield high performance. In imitating complex artefacts, a small error in understanding can have large repercussions, as the economic value of complex artefacts lies precisely in the complementarities between its parts (Rivkin, 2000). While learning is especially difficult for firms dealing with complex products, innovation in complex product industries nevertheless relies a lot on inter-firm learning (Miller et al., 1995). Firms going alone by relying fully on their internal knowledge are bound to end up in poor local optima. Hence, though difficult, firms in complex product environments benefit from supplementing their internal innovation efforts (exploitation) with learning by imitating knowledge from others (exploration).

Errors are also more likely to occur, the more two firms differ in their knowledge (Nooteboom, 2000). Indeed, to understand and use new knowledge from others, a certain level of absorptive capacity is required. Organizations will find it much easier to learn from firms that have much knowledge in common, as the new knowledge learnt will be more easily understood and combined with the existing knowledge base (Cohen & Levinthal, 1990; Solís-Molina, Hernández-Espallardo, & Rodríguez-Orejuela, 2018). Yet, as theorized by Nooteboom, a fundamental trade-off is implied in inter-organization learning *"between cognitive distance, for the sake of novelty, and cognitive proximity, for the sake of efficient absorption. Information is useless if it is not new, but it is also useless if it is so new that it cannot be understood"* (Nooteboom, 2000, p. 72). The trade-off captures the two sides of exploration: on the one hand a firm seeks to learn from others exactly by exploring very new knowledge suggesting it should look for cognitive distant firms, while on the other hand a firm wants to avoid copying errors by looking at cognitive proximate firms from which it can easily absorb knowledge. Hence, one expects that there exists an optimal cognitive distance between two firms that maximizes the benefits of learning by one firm from the other firm (Nooteboom, 2000, p. 74).

While the concept of absorptive capacity emphasizes the cognitive differences between firms, the extent to which firms can gain access to knowledge held by firms also depends on social contacts (Uzzi, 1996). Many employees of firms maintain social ties with employees in competing firms, and use such relations for informal knowledge sharing (Bouty, 2000). What is more, mutual sharing practices are reinforced by professional and academic norms in communities of practices (Lissoni, 2001). Acquaintanceships may stem from having been colleagues in the past, having been fellow students in the same school or university, or having been collaborators in joint projects (including past license agreements and strategic alliances). Employees engage in informal knowledge sharing as it raises their own expertise despite a possible loss of competitive advantage of the imitated firm. Such losses may anyway be small, as employees who share knowledge unilaterally do so with the expectation that the favor will be returned at a later moment in time (Bouty, 2000).

Informal knowledge sharing among employees supports imitation processes between firms (Uzzi, 1996). Given the importance of networks as a source of knowledge spillovers, the position of firms within networks channeling knowledge spillovers will thus affect its ability to learn and to innovate (Breschi & Lissoni, 2009; Powell et al., 1996; Pyka, 2002). Indeed, there is ample empirical evidence that the characteristics of firm networks through which knowledge flows take place, and the position of firms within such networks, are relevant to firm performance (for reviews: Özman, 2009; Phelps, Heidl, & Wadhwa, 2012).

Past research highlighted two distinct network characteristics as relevant to inter-organizational learning. First, learning depends on the extent to which an organization has access to knowledge held by others. From a network perspective, however, access does not only depend on an organization's direct ties ('friends' at social distance 1), but also on its indirect ties ('friends of friends' at social distance larger than 1) (Ahuja, 2000; Breschi & Lissoni, 2009). In general, one can expect that knowledge from socially proximate firms is more accessible than knowledge from socially distant firms. This has been confirmed by empirical research showing that inter-firm patent citations occur less often, the more distant two firms are in the social network of co-inventors (Breschi & Lissoni, 2009). Hence, regarding access to knowledge, the value of organization's network position can be expressed by the average social proximity to all other firms.

Second, it has been argued that network clustering in triangle relationships matters ('friends of friends being friends'). Triangle relationships support trust as actors have fewer incentives to behave opportunistically, as opportunistic behavior towards one partner may jeopardize the relation with the other partner in a triangle (Granovetter, 1985). Trust, in turn, supports the exchange of valuable knowledge and collaborative problem-solving (Uzzi, 1996). Thus, while short distances provide access to a wide range of different ideas, clustering provides a complementary structure allowing organizations to elaborate upon selected ideas in close collaboration. That is, short distances and clustering in networks are likely to be complements: short distances are associated with the exploration of new ideas and clustering with the further detailed elaboration of such ideas (Capaldo, 2007; Fleming, Chen, et al., 2007; Schilling & Phelps, 2007).

Another strand of literature, mainly in the field of economic geography, has further unpacked different forms of proximity. In particular, it has been argued that cognitive and social proximity may act as substitutes (Boschma, 2005; Huber, 2012). If cognitive proximity between firms is high, social networks may not be required for effective transfer to take place, as the imitating firm can easily grasp and 'reverse engineer' the solution of its competitor. If cognitive proximity is low, by contrast, social networks may be crucial for the imitating firm as its employees have informal access to knowledge residing in the other firm. Hence, next to the positive effect of cognitive proximity and social proximity on imitation, one can further hypothesize that high social proximity is especially supportive of imitation in context of low cognitive proximity and, *vice versa*, that low social proximity is sufficient for imitation in context of high cognitive proximity (Huber, 2012).

To combine the absorptive capacity and social network arguments in a single exploration-exploitation framework, we will propose a model in which firms, cyclically, engage in local search (exploitation) and then imitate competitors with higher performance (exploration), and so on. All firms are assumed to be part of a single social network. The network's structure, stemming from the informal social networks maintained by their employees outside the control of the firm, is exogenously given. We investigate the effect of network structure on the average performance of firms, by comparing networks that differ in terms of the average distance between firms and the average clustering between firms using the 'small-world' parameter (further explained below). We also analyze dyads (pairs of firms) to see if the model can replicate the empirical findings that successful imitations typically occur at an intermediate cognitive distance and that (rarer) successful imitations between cognitively distant firms require high social proximity.

3. The model

To investigate the role of social networks in imitation efforts among firms, we use a simulation model starting from the NK-model of fitness landscapes (Levinthal, 1997) in which firms innovate while being part of a small-world network (Watts & Strogatz, 1998). A simulation model allows one to systematically evaluate the effect of exogenous parameters on the individual firm-level and collective industry-level performance, which are the foci of management and economics scholars, respectively. The key parameters here are the complexity of the problem at hand (as expressed by K in the NK-model) and the degree of randomness in the small-world network (as expressed by β in the small-world model).

We approach the innovation logic of exploration and exploitation as follows. Each firm performs innovation on an object of a certain complexity, with object and complexity being the same for all firms. We thus consider the complexity of a firm's product or service as exogenous to a firm, while specific to industries. For example, aerospace, automobile and information technology industries are generally considered complex product industries, where the final product consists of many different parts and underlying knowledge bases, while products such as furniture, toys and clothing can be considered to be less complex products (Marsili, 2002). The different levels of complexity as expressed by parameter K can thus be understood as representing different industry contexts.

We assume that firms are all engaging in both exploration and

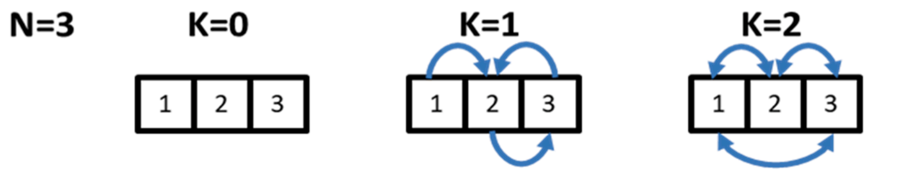

123	p1	p2	p3	P	p1	p2	p3	P	p1	p2	p3	P
000	0.6	0.3	0.5	0.47	0.6	0.3	0.5	0.47	0.6	0.3	0.5	0.47
001	0.6	0.3	0.9	0.60	0.6	0.5	0.9	0.67	0.1	0.5	0.9	0.50
010	0.6	0.8	0.5	0.63	0.4	0.8	0.1	0.43	0.4	0.8	0.1	0.43
011	0.6	0.8	0.9	0.77	0.4	0.5	0.8	0.57	0.3	0.5	0.8	0.53
100	0.9	0.3	0.5	0.57	0.9	0.3	0.5	0.57	0.9	0.9	0.7	0.83
101	0.9	0.3	0.9	0.70	0.9	0.5	0.9	0.77	0.7	0.2	0.3	0.40
110	0.9	0.8	0.5	0.73	0.6	0.8	0.1	0.50	0.6	0.7	0.6	0.63
111	0.9	0.8	0.9	0.87	0.6	0.5	0.8	0.63	0.7	0.9	0.5	0.70

Fig. 1. Examples of performance calculations for K = 0, K = 1 and K = 2 (N = 3). For K = 0, no interdependencies between components exist, so each component has only two performance values – one for each of its own states. For K = 1, in this example, component 1 depends on component 2, component 2 on component 3, and component 3 on component 2. For each component, there are four possible performance values. For K = 2, all components depend on all other components, resulting in eight possible performance values for each component. The overall performance *P* of each string is the average of the performance values of the component states (Kauffman, 1993).

exploitation, yet in a temporal order (Gupta et al., 2006). Each firm starts innovating by internal local search thus exploiting its existing knowledge base. Once exploitation reaches a local optimum, the innovation process turns to exploration in an attempt to escape the local optimum through a 'long jump' (Levinthal, 1997). In our model, we approach exploration as gaining knowledge from other firms by imitation (Csaszar & Siggelkow, 2010). Once incorporated, this new knowledge provides a basis for another cycle of exploitation, which again will end once a new local optimum is reached, triggering exploration by imitation again, and so on. A firm engaging in exploration imitates another firm, where the imitated firm is selected on the basis of its relative performance (Nelson & Winter, 1982), while imitation is error-prone (Rivkin, 2000). We assume that such errors are more likely to occur if two firms are socially more distant. The social distance between each pair of firms is derived as the shortest path ('geodesic distance') between two firms in the social network at hand. To explain the model, we first describe below how a focal firm exploits its knowledge internally until it reaches a local optimum (exploitation) and then how this firm acquires new knowledge by imitating another firm (exploration) with a certain degree of fidelity depending on the social distance between the imitating firm and the imitated firm.

3.1. Exploitation

Exploitation is modelled here using the NK-model (Kauffman, 1993). This model was originally developed in the context of biology for the study of interdependence between genes in a genome, known as epistasis. Epistatic structures are not confined to biological systems, but are also typical for technological components making up a technology and organizational tasks making up a production process (Levinthal, 1997). By now, the NK-model has become a generic model of search in the management literature as systematically reviewed by Granco and Hoetker (2009), Puranam, Stieglitz, Osman, and Pillutla (2015) and Baumann et al. (2019).

In the NK-model, a firm's knowledge base is represented as a string of N components. The level of complexity faced by a firm depends on the number of interdependencies between components modelled by parameter K and ranging from $K = 0$ to $K = N - 1$. Complexity here implies that the performance, or *fitness*, of each component depends on the state of K other components. The performance P of a string is given by the average performance over all components, with the performance value of each component depending on the state of the component in question and the K components it depends on. Without loss of generality, components can be in two states: either 1 or 0. Performance values, randomly drawn from a uniform distribution [0, 1], are assigned to each of the 2^{K+1} unique combinations of the state of a component and the states of the K components it depends on. With $K = 0$, this means that for each component, performance values are only drawn for each of its two possible states. With $K = 1$, this means that for each component, performance values are randomly drawn for each of the two possible states of that component combined with each of the two possible states of the component it is depending on, and so on (see, for an example, Fig. 1).

With N components, each having two possible states, there are 2^N possible configurations. Each configuration is represented in our model as a bitstring. The difference between any two configurations is expressed by the number of components with a different state. This difference is called the Hamming distance H between two bitstrings (Hamming, 1950). For example, for bitstring 01011 and bitstring 00011, we have H = 1, while for bitstring 01011 and bitstring 10100, we have H = 5. In our model, the Hamming distance H between two bitstrings captures the cognitive distance between two firms, with its reverse (N – H) denoting the cognitive proximity between firms.

The 2^N possible configurations of bitstrings can be used to construct a 'landscape' (Wright, 1932), with the height of each bitstring given by its performance and the distance between bitstrings in the landscape by the Hamming distance. A peak in the landscape is a bitstring that has a superior performance compared to all its N neighbors at a Hamming distance equal to 1. The highest peak corresponds to the configuration with the best performance, which is the global optimum, while other peaks correspond to configurations that are local optima. For example, in Fig. 1, for $K = 2$, bitstring 100 is the global optimum and bitstring 111 is a local optimum. It follows that the minimum Hamming distance between two optima is two.

Assuming that firms exploit their knowledge base through local

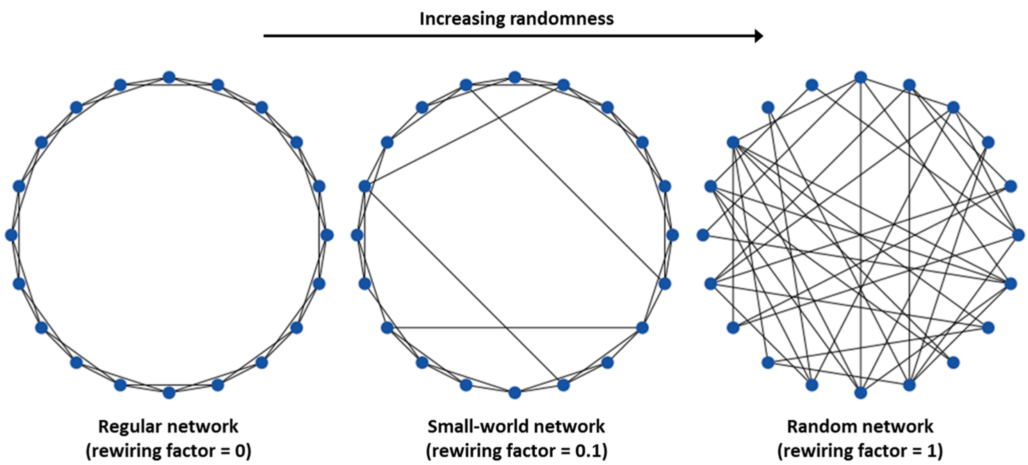

Fig. 2. Three examples of networks of firms (F = 20, D = 4). Regular networks have high clustering. Networks are created by following the method of Watts and Strogatz (1998), varying the 'rewiring factor' β between 0 and 1.

search by successively moving to neighboring bitstrings with higher fitness, a firm is bound to end up at a local optimum, where it ends its exploitation activity. The key insight derived from the NK-model holds that, for a given N, the number of local optima increases with complexity K. This implies that the more complex the knowledge base of a firm, the more likely it will get quickly 'stuck' on a suboptimal performance level (Levinthal, 1997). Exploitation, then, will only be effective in finding a local optimum and subsequent exploration is needed for search to continue. Exploration, then, can be metaphorically thought of as a 'long jump' away from a local optimum (Levinthal, 1997).

3.2. Exploration

In our model, exploration activities concern efforts of firms to imitate better performing firms. We thus assume that firms cannot protect themselves from imitation by others. One may think of imitation as reverse engineering benefitting only the imitating firm. Alternatively, imitation can be thought of as following from a transaction between the two firms (e.g., through licensing), which would decrease the current profits of the imitating firms and would increase the profits of the imitated firm. Hence, the performance values in the model refer to the value of the knowledge base (e.g., technical efficiency of a technology) and not to the economic value of the firm as such (e.g., the profitability of a firm). The exact mechanism underlying imitation, however, does not bear any implications for the model setup.

In the population of F firms, we thus assume that each firm can observe the performance P of all other $F - 1$ firms. In its exploration activities, a firm focuses its attention only to the subset of F' better performing firms f_k. A firm engaging in exploration imitates another firm, where the imitated firm is selected on the basis of its relative performance (Nelson & Winter, 1982). In an act of exploration, then, a firm f_i imitates another firm f_j with a probability π_{ij} proportional to the latter's relative performance among all F' better performing firms:

$$\pi_{ij} = \frac{P(f_j)}{\sum_{k=1}^{F'} P(f_k)}$$

also known as 'roulette wheel selection' (Goldberg, 1989).

If firm f_i decides to imitate firm f_j, firm f_i will attempt to copy each bit from the string occupied by firm f_j that is different from its own. Under the assumption of perfect imitation, the imitating firm f_i will simply substitute all the H component states in its bitstring that are different from the bitstring of imitated firm f_j, and consequently the imitating firm will attain the same performance level as the imitated firm This means that the higher the cognitive distance between two firms (that is, the Hamming distance H between their bitstrings), the more the imitating firm will learn from the imitated firm, the longer its 'jump' in the fitness landscape will be.

In a setting of perfect imitation, the outcome of our model of exploitation and exploration will be rather unsurprising. If all firms start from a randomly assigned string, they will first engage in exploitation by local search until they reach a local peak, and then start imitating better performing firms by exploration. It follows that the best performing firm will not engage in imitation and that the second best performing firm will always aim to imitate the best performing firm. It also follows that in the absence of copying mistakes, all firms will eventually converge to the highest local optimum found by some firm in the first stage of exploitation. Note that the chance this local optimum is also the global optimum will increase with the number of firms in the population F and will decrease with the complexity K.

The assumption, however, that firms can imitate without errors is a very strong one. As argued before, firms generally make mistakes when imitating other firms due to their limited absorptive capacity (Cohen & Levinthal, 1990; Rivkin, 2000). This means that imitation here is the *attempt* by firm f_i to copy all the H bits from the bitstring of firm f_j that are different from f_i's own bitstring. However, while attempting to perfectly copy the H bits of the other firm, it may makes mistakes in doing so. Hence, given a certain probability of any copying mistake occurring, it follows that the higher the cognitive (Hamming) distance between two firms is, the more likely that at least one copying mistake will occur. This thus captures the idea that cognitive distance between firms renders imitation less effective (Nooteboom, 2000).

The complexity of the knowledge imitated also matters (Rivkin, 2000). If there exist few interdependencies between components (low K), one copying error affects only few other components, while if there are many interdependencies between components (high K), a single

Table 1
Parameter settings and output variables.

Parameter	Description	Settings
F	Number of firms	100
D	Degree	6
β	Rewiring factor	0, 0.01, 0.1, 1
N	Number of bits	16
K	Complexity	3, 7, 11, 15
Output variable	**Description**	**Range**
P	Performance ('fitness')	[0, 1]
Cognitive distance	Hamming distance H	[1, N]
Social distance	Shortest path δ	[1, Δ]

copying mistake will affect many other components. In the extreme case of maximum complexity ($K = N - 1$), a single copying mistake will lead a firm to discover a string with a performance level uncorrelated to the performance level of the string it attempted to copy, because, in the case of maximum complexity, the fitness values of all components are redrawn moving from one string to a neighboring string. The expected fitness value, then, of the wrong string discovered is then simply the expected value of any random draw, being 0.5.

Rather than assuming that copying mistakes occur with some constant positive probability, we assume that the social distance between the imitating and the imitated firm affects the probability of errors occurring in imitation. Personal ties between firms, be them direct ('friends') or indirect ('friends of friends') ties, are assumed to support informal knowledge sharing and to enhance the accuracy of imitation (Bouty, 2000; Breschi & Lissoni, 2009). The lower the social distance between two firms, then, the less likely the imitating firm will make mistakes in copying a bit from the string the imitated firm.

The social distance between any two firms in the social network is defined by the shortest path length between firm f_i and firm f_j indicated by δ_{ij}. The probability of an error occurring in copying a bit, then, is given by the shortest distance between the two firms divided by the shortest path between the two most distant firms in the network Δ (known as the network's diameter). Thus, each bit that firm f_i attempts to imitate from firm f_j will be copied with error with probability $\frac{\delta_{ij}}{\Delta}$. It follows that neighboring firms make such errors only at a rate of $\frac{1}{\Delta}$, while two firms at the maximum distance of Δ will be unable to imitate each other, as all the bits will be copied with error. It also follows that for a firm to make a perfect copy, it should avoid to make any mistake in the copying of H bits. One can thus express the probability of a perfect copy as $(1 - \frac{\delta_{ij}}{\Delta})^H$.

Note that with the introduction of social proximity in the model, the detrimental role of cognitive distance still remains intact. For any value of social proximity, the probability of copying mistakes applies to any of the bits that a firm tries to copy. Hence, the number of copying mistakes in imitation still remains proportional to the Hamming distance between the bitstring of the imitating firm and the bitstring of the imitated firm.

We consider four different types of networks ranging from a perfectly regular network, to two 'small-world networks', to a completely random network. We look at sparse networks where the number of ties of each of the firms (called 'degree' D) is rather small compared to the number of firms F. We start from a regular network arranging firms on a circle with each firm having the same degree and the same average distance to all other firms. Starting from this regular network, irregularity is introduced by randomly rewiring a certain fraction β of links between firms, with rewiring factor β tuned between $\beta = 0$ (fully regular network), to $\beta = 0.01$ and $\beta = 0.1$ (small-world), to $\beta = 1$ (fully random network) (Watts & Strogatz, 1998).

As demonstrated by Watts and Strogatz (1998), the rewiring of links in a regular network has a profound effect on the average path length between every two firms (i.e., the average social distance). Small-world networks, in particular those with a rewiring factor between $\beta = 0.01$ and $\beta = 0.1$, maintain the high degree of clustering characteristic of regular networks, but have a much shorter average path length than regular networks, as the small fraction of rewired ties function as 'shortcuts' (Fig. 2). When further rewiring all ties as to obtain a random network, the clustering of nods in triangles gets lost while distance between nodes become even shorter. Thus, small-world networks combine the feature of high clustering from regular networks and the feature of short distances from random networks.

There is extensive empirical evidence that social networks between acquaintances and between firms exhibit the small-world features of short distances and high clustering (Fleming, King, & Juda, 2007; Lissoni, Llerena, & Sanditov, 2013; Newman, 2003; Uzzi, Amaral, & Reed-Tsochas, 2007). Given the empirical relevance of small-world networks, we will focus in the simulation results primarily on the intermediate values of the rewiring parameters, while considering the minimum and maximum value of theoretical relevance, mainly.

3.3. Simulation

In each simulation, firms start from randomly assigned bitstrings and engage in search for alternating periods of exploitation and exploration. Exploitation consists of local search until a local peak is reached. Once a local peak is reached, a firm tries to imitate a better-performing firm. If exploration is successful, a firm reverts again to exploitation, and the cycle repeats.

In our model, exploitation implies local search which means a one-bit mutation (from 0 to 1 or from 1 to 0). We apply an algorithm that is known as 'greedy search' (Goldberg, 1989), meaning that a firm in each hill climbing step will mutate the component that will generate the best performance gain when changed. In other words, a firm takes the steepest way up from its current position in the landscape, hoping to follow the shortest route to a nearby peak. Greedy search steps are repeated until no more improvements can be found.

Once a firm has reached a local optimum, only exploration may yield higher fitness. Imitation attempts will only be accepted if fitness increases. Yet, as imitation is failure prone, an imperfect imitation may well result in finding a bitstring with a lower fitness than the firm's current local optimum. In that case, a firm remains at its current position and attempts again in the next time-step. Also note that in some cases, imperfect imitation may actually lead to a higher fitness that initially envisaged (Posen, Lee, & Yi, 2013). In that case, the firm will accept the newly found bitstring as the fitness of the new string is higher than the previous string.

In each simulation, firms are allowed to innovate in a series of greedy search and imitation steps, limited to a maximum of 200 steps.[1] Based on the model as described above, we have carried out a series of simulations with parameters values and output variables as given in Table 1. We set the parameters values for the number of firms (F = 100), their degree (D = 6) and the length of the bitstring (N = 16). What we vary is the small-world rewiring factor β as to compare different types of network and the parameter K to compare different levels of complexity. To prepare our simulations, we have generated a set of 100 networks for each rewiring factor $\beta > 0$ (for $\beta = 0$, the network is a given). From these 100 networks, we selected ten representative networks with modal values of average clustering and average path length. Likewise, we have prepared a set of ten NK landscapes for each K-value and selected the most representative landscape with the modal number of optima.

The model was implemented in NetLogo (Wilensky, 1999). For each simulation, with 100 runs for each combination of K and rewiring factor β, the model was executed according to the following pseudo code:

[1] The maximum number of time steps was determined experimentally by a series of trial simulations, which demonstrated that at this time step, firms have reached a stage in which performance gain levels have become minimal.

```
;initiation
random-select-network (β parameter, F=100)
Δ = network-diameter
random-select-NK-landscape (K parameter, N=16)
for firm f=1 to F
      bitstring (f) = random-bitstring (N=16)
next firm f
for time-step = 1 to 200
      ;exploitation
      for firm f=1 to F
            repeat until local-optimum = found
                  bitstring (f) = exploitation-by-greedy-search (f)
            end repeat
            bitstring' (f) = bitstring (f)
      next firm f
      ;exploration
      for firm f=1 to F
            f' = imitation-firm-by-roulette-wheel-selection (f)
            sd = social-distance (shortest-network-path f f')
            cd = cognitive-distance (hamming-distance f f')
            b-current = bitstring (f)
            b-new = imitate-with-error (bitstring'(f'),  sd, Δ, cd)
            if performance(b-new) > performance(b-current) then
                  bitstring(f) = b-new
            end if
      next firm f
next time-step
```

4. Results

4.1. Network-level

Fig. 3 provides the time evolution of the average performance of firms in the population for different network types and complexity levels. We express the performance level of firms as the percentage increase in fitness compared to the average fitness of a simulation run without exploration. The percentage expresses how much the average fitness of local optima found when search includes both exploitation and exploration, exceeds the average fitness of local optima found if firms only engage in exploitation, that is, a single round of greedy search. These results allow us to understand the value of exploration on top of exploitation for different types of networks and complexity levels.

The first finding that can be derived from Fig. 3 holds that, indeed, exploration contributes to firm performance, as it allows a firm to escape a poor local optimum through imitation. The extent to which exploration helps to improve firm performance increases with the complexity of the knowledge base. This can be understood from the fact that less complex landscapes have fewer local optima among which the best optima have the largest basins of attraction (Kauffman, 1993). Exploitation alone, then, will often lead firms to optima with reasonably high fitness. More complex landscapes, by contrast, have many more local optima with only small basins of attraction. Exploitation alone, then, will often lead firms to poor local optima. The differences are quite pronounced: exploration adds some six to seven percent to firm performance for $K = 3$, and some 16 percent to firm performance at $K = 15$.[2] This finding is in line with empirical research based on patent data showing that the higher the level of technological interdependence in an industry, the more important exploration activities are to improve firms' performance (Gatti et al., 2015).

The second finding that one can derive from Fig. 3 holds that small-world networks outperform regular and random networks only in complex landscapes. Recall that small-world networks combine the characteristic of high clustering from regular networks and the characteristic of short distances from random networks. The two properties have different effects on exploration. Short distances have the obvious effect that imitation becomes more effective as fewer errors are made between socially proximate firms. Clustering, however, has a subtler effect. On the one hand, clustering leads firms to converge faster on a local optimum as learning takes place in triangles. At the same time, as the social distance between firms in more clustered networks is higher compared to less clustered networks, clustering inhibits imitation between firms in different clusters, thus reducing premature convergence at the population level (Baumann et al., 2019). Put differently, clustering is advantageous to maintain a certain variety in knowledge to fuel future innovation, a logic which has been highlighted as a generic principle in cultural evolution (Muthukrishna & Henrich, 2016).

Looking closer at the results in Fig. 3, we see that for low complexity ($K = 3$) all networks perform equally well as firms perfectly converge in fitness levels (due to the consistent discovery of the global optimum in all simulation runs). For moderate complexity, the random networks perform best although differences are small. Once complexity becomes even higher, small-worlds start to outperform the other networks. In complex landscapes, the number of local optima is high leading to a sustained diversity of bitstrings. Clustering then, turns into a positive force helping firms in cliques to filter out high optima without leading to premature convergence ideas. This is further evident from comparing the two small-world networks, where the one with the highest clustering ($\beta = 0.01$) outperforms the one with lower clustering ($\beta = 0.1$) for the highest level of complexity ($K = 15$).

4.2. Dyad-level

Fig. 4 provides a heat map representing the frequencies at which successful imitations occurred between any two firms at a certain social and cognitive proximity. We show a total of sixteen of such maps corresponding to the parameter space given by the four different network structures (β) and the four different complexity levels (K).

As a preliminary observation, we find that imitation is seldom successful if firms are distant in both the social and the cognitive dimension. This comes naturally out of the model as cognitive distance increases the number of bits that a firm tries to copy and social distance increases the probability of a copying mistake in any bit. A second observation holds that for more random networks, imitation occurs at lower social proximity. This reflects the short distances in random networks.

As a first substantial finding, we observe that successful imitation happens most often at an intermediate level of cognitive proximity. As expected, imitation at high cognitive distance is hardly feasible given

[2] Comparing the means of the final values (at time step 200) of simulations with different rewiring factors (β), we find that the differences between means were significant ($p < 0.05$), except for simulations for $K = 3$.

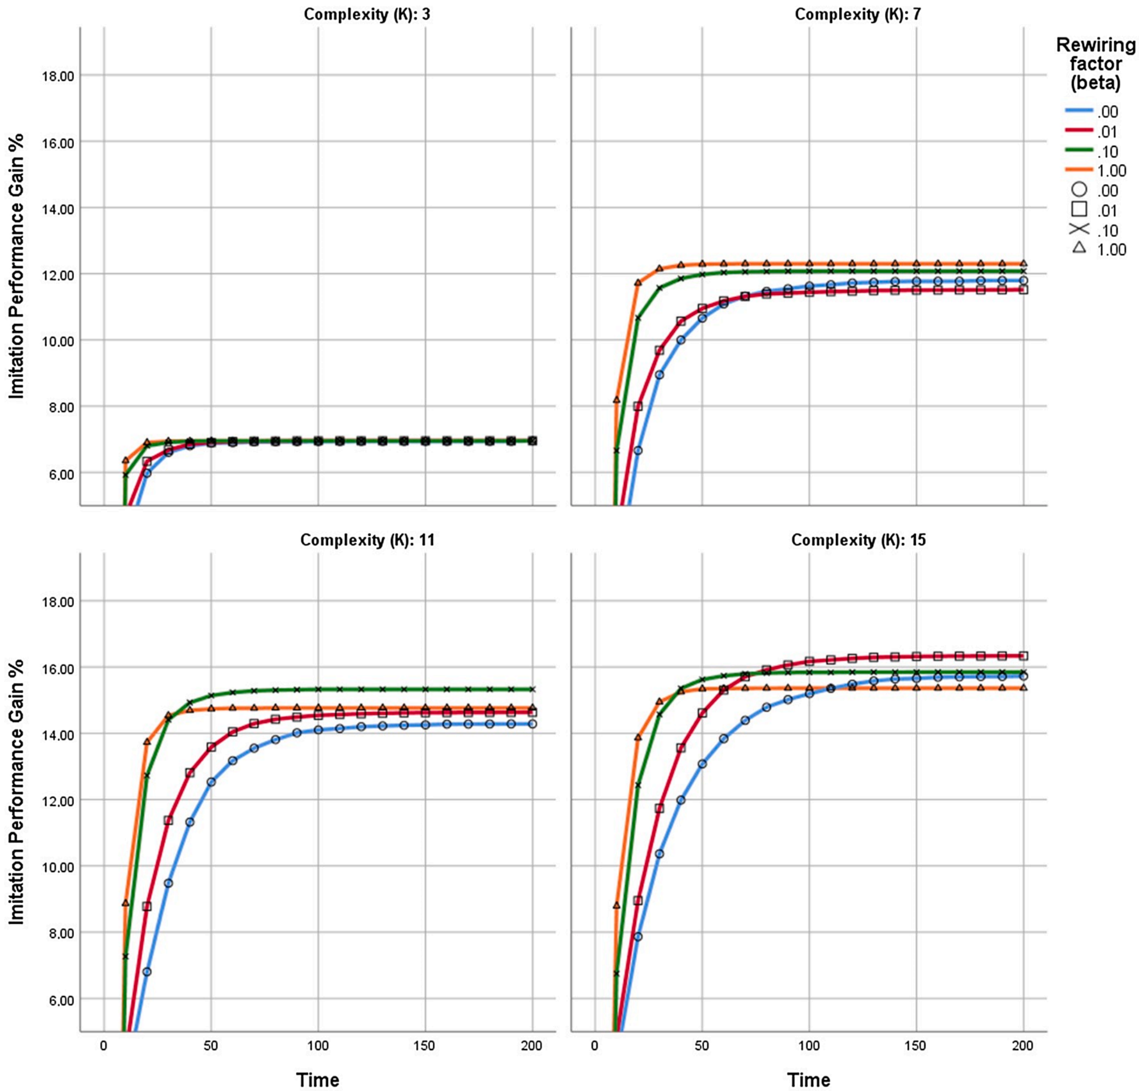

Fig. 3. Performance gain achieved by exploration compared to the baseline with firms only engaging in exploitation. Averages over 100 simulation runs.

the higher chances of copying errors. Imitation at low cognitive distance, by contrast, should be most effective in that it is likely to occur without mistakes. However, imitation of similar bitstrings rarely happens as most nearby bitstrings will have inferior performance. Hence, we can understand that successful imitation occurs most often at intermediate distance. Interestingly, this result holds regardless of the type of network considered. The optimal cognitive proximity pattern also holds for different complexity levels, except for the case of lowest complexity ($K = 3$) where no pronounced pattern is observable. As most of the performance gains are achieved by exploitation rather than by exploration in a low-complexity context as we just discussed, the inverted-U pattern is not well discernable. For higher complexity levels the inverted-U pattern is robust. Our simulation results displaying an optimal cognitive proximity are thus consistent with the results coming out of empirical studies (Fitjar, Huber, & Rodríguez-Pose, 2016; Gilsing, Nooteboom, Van Haverbeke, Duysters, & Van den Oord, 2008; Nooteboom, Van Haverbeke, Duysters, Gilsing, & Van den Oord, 2007; Wuyts, Colombo, Dutta, & Nooteboom, 2005).

As a second substantial finding, we see that social and cognitive proximity act as substitutes, a thesis advanced earlier by proximity researchers (Boschma, 2005; Huber, 2012). We find that successful imitation is more common among cognitive distant firms if their social proximity is high, while reversely, successful imitation is also more common among socially distant firms if their cognitive proximity is high. This shows that longer jumps in the landscape are supported by social proximity as the chances of copying mistakes go down and, accordingly, the chances of finding a higher fitness go up. The results also show that firms can learn from unfamiliar competitors at a large social distance, as long as these firms work in the same knowledge area. As a further qualification, it is clear that the inverse relation between social and cognitive proximity is less apparent in random networks where social distances are small anyway. The model results are consistent with the empirical studies that found social proximity is especially important when cognitive proximity is low, and *vice versa*, that cognitive proximity is especially important when social proximity is low (Cassi & Plunket, 2015; Huber, 2012; Steinmo & Rasmussen, 2016).

4.3. Sensitivity analysis

In the presentation of results, we focused on differences between network structures (β) and complexity (K) while holding the other parameters constant. This leaves open the question whether our results are robust when varying the other three parameters (F, N, and D).

When we increase (decrease) the number of firms F in our network, the mean performance achieved by firms will improve (deteriorate). This result can be simply explained by the fact that once more (fewer) firms are active in the same landscape, the probability of a finding a high peak will increase (decrease).

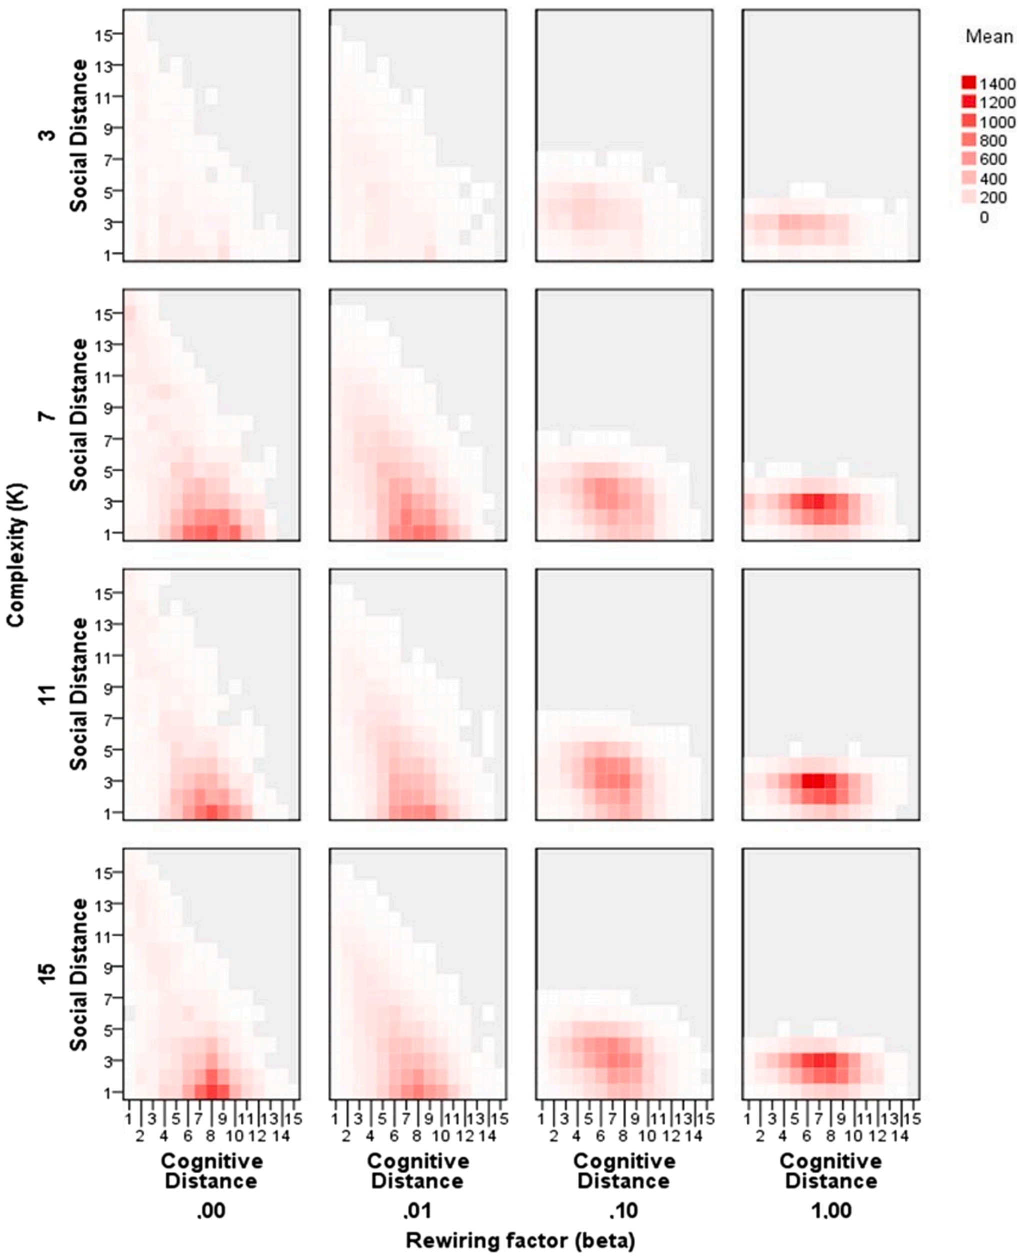

Fig. 4. Occurrence (darker color corresponds with higher occurrence) of successful imitation, for different complexity (K) levels (rows) and rewiring factors (β) (columns), at different cognitive and social distances (reverse of proximity). Means taken over 100 simulation runs. For lower complexity (K = 3) imitation is less prominent, especially in more clustered networks. In case of more complexity, successful imitation is most common at intermediate cognitive distance (inverted U-shape). For social proximity, the expected U-shape is most prominent in less clustered or more random networks, but absent in highly clustered networks.

When increasing (decreasing) N while adapting the K-values accordingly (keeping K-values proportional to N as in the original simulation), imitation will be more difficult as the Hamming distances between the imitating firm and the imitated firm go up. This implies that the number of errors goes up (goes down) and the learning rate goes down (goes up).

Finally, increasing (decreasing) the degree D in the social network will lower the average social proximity between firms leading to fewer (more) errors in exploration activity. As a result, firms will learn more (less) effectively and will find local optima with high (lower) fitness.

5. Conclusion

In our theoretical model, we consider innovation as a process alternating between internal exploitation by local search and external exploration by imitation of successful others. In the model, multiple firms search for optima in an NK fitness landscape (Kauffman, 1993) while being connected in a small-world social network (Watts & Strogatz, 1998). A firm's network position is assumed to affect the fidelity of its imitation efforts, with the probability of mistakes increasing with the social distance between the imitating firm and the imitated firm.

The key result of the model holds that successful (i.e. fitness increasing) imitation typically occurs at an intermediate level of cognitive proximity, consistent with empirical studies. At low cognitive distance there are rarely successful firms to imitate, while at high cognitive distance imitation often fails due to copying errors. The second key result holds that social and cognitive proximity are substitutes, also found in empirical studies. Successful imitation is more common among cognitively distant firms if their social proximity is high, while reversely, successful imitation is also more common among socially distant firms if their cognitive proximity is high. Apart from these two stylized facts, the model reproduces two more stylized facts: the higher value of exploration in highly complex industries compared to less complex industries, and the benefits of the small-world network structure for collective learning compared to regular and random network structures.

The main theoretical implication of our model concerns the key role of social networks among firms in supporting effective imitation. A firm with a central network position has short social distances to other firms allowing it to imitate effectively a large set of firms with varying knowledge bases, which is shown to be of particular relevance in complex product industries. By comparing different network structures at the population level and learning at the level of two firms, we further have been able to integrate the theory of small-worlds in collective learning and the proximity theory regarding inter-firm learning.

Our model exemplifies the usefulness of the NK-model as a framework to investigate exploitation as local search to a peak and exploration as a jump away from a peak (Csaszar & Siggelkow, 2010). We choose to model exploration as consisting only of imitation activities. The model does not address learning within strategic alliances that are set-up for mutual knowledge exchange. Learning within such alliances is arguably quite distinct from imitation, as alliances are cooperative structures aiming to generate new knowledge for both partners through recombination ('crossover') rather than through imitation (Cowan & Jonard, 2009). To investigate the generalizability of our results, a future model may systematically compare the outcomes of exploration via imitation versus exploration via alliances.

Our model can also be extended in other ways. First, apart from exploration by imitation as we do here, the model may incorporate exploration by internal search as in the original NK-model (Levinthal, 1997). Then, a question to address holds what optimal balance exists between internal and external exploration depending on a firm's competitive position. In particular, well-performing firms can arguably learn more from exploration by internal search compared to poorly performing firms that can benefit more from imitating others. A second extension is to incorporate cost. In our model, we abstracted away from the cost of imitation. While we capture the higher probability of errors when imitation distance goes up, one could further argue that copying more bits does not only entail more risk, but also higher costs (Csaszar & Siggelkow, 2010). A final extension regards the investigation of environmental turbulence on exploration and exploitation, which can be integrated with the NK-model by making fitness level noisy (Uotila, 2018). Here, a key question holds whether imitation is still as effective as a means to conduct exploration if fitness levels of competitors convey noisy information.

The managerial implications that follow, in its most general sense, are threefold. First, firms profit from participating in small-world networks in contexts of high product complexity. Hence, strategically maneuvering into a favorable network position combining short distances with high clustering only matters in contexts where products, and the underlying knowledge base, are complex. Second, firms should refrain from attempting to imitate firms with a very similar knowledge base as well as from firms with a very dissimilar knowledge base. Instead, firms should focus their learning efforts at a particular subset of fellow firms that are sufficiently different to effectively learn from, but not too different as to avoid the risk of failure in learning. Third, firms who specifically aim to learn distant knowledge should invest in social proximity. This can be done by encouraging labor mobility of their own employees or poaching employees from competitors.

Along similar lines, some general implications for government policy can be derived. First, in contexts of high knowledge complexity, governments can try to influence inter-firm network structures in ways that the overall network structure acquires small-world characteristics. As a public actor, government can influence the macro-level structures of collaboration among firms by subsidizing inter-firm networks (Van Rijnsoever, Van Den Berg, Koch, & Hekkert, 2015). In particular, governments can focus on creating shortcuts between two socially distant firms as to increase social proximity, and on promoting large consortia among firms in contexts where the level of clustering is low as to increase clustering. In addition, governments can promote social proximity more generally if it wants to promote learning across unrelated domains. For example, promoting associational life in general, and 'policy platforms' and 'innovation intermediaries' in particular, are ways to bring together businesses in an open setting (Janssen & Frenken, 2019).

As our model is an abstract one, the implications that can be derived from the model may stretch beyond the immediate context in which we presented it. In particular, the notion of social proximity can be extended to any form of proximity that affects the fidelity of learning, including geographical proximity supportive of face-to-face interactions as often happens in geographic clusters (Boschma, 2005). Note here that geographical proximity also exhibits the small-world network logic of our model in that local interactions can be associated with high trust supportive of collaborative elaboration of existing ideas and the global interactions with the short cuts needed to bring in new knowledge from abroad. *Mutatis mutandis*, the policy implications for firms and governments would hold that geographical proximity among firms, organized in 'industrial clusters', is especially relevant in industry contexts with high knowledge complexity.

Acknowledgements

We thank Floor Alkemade (Eindhoven University of Technology) and the anonymous reviewers for useful feedback. The usual caveat applies.

Funding

Koen Frenken received funding from the Netherlands Organisation for Scientific Research (NWO) within the Vici-program (number 453-14-014).

References

Ahuja, G. (2000). Collaboration networks, structural holes, and innovation. A longitudinal study. *Administrative Science Quarterly, 45*, 425–455.
Baumann, O., Schmidt, J., & Stieglitz, J. (2019). Effective search in rugged performance landscapes: A review and outlook. *Journal of Management, 45*, 285–318.
Benner, M. J., & Tushman, M. L. (2003). Exploitation, exploration, and process management: The productivity dilemma revisited. *Academy of Management Review, 28*, 238–256.
Boschma, R. (2005). Proximity and innovation: A critical assessment. *Regional Studies, 39*, 61–74.
Bouty, I. (2000). Interpersonal and interaction influences on informal resource exchanges between R&D researches across organizational boundaries. *Academy of Management Journal, 43*, 50–65.
Breschi, S., & Lissoni, F. (2009). Mobility of inventors and networks of collaboration: An anatomy of localised knowledge flows. *Journal of Economic Geography, 9*, 439–468.
Capaldo, A. (2007). Network structure and innovation: The leveraging of a dual network as a distinctive relational capability. *Strategic Management Journal, 28*, 585–608.
Cassi, L., & Plunket, A. (2015). Research collaboration in co-inventor networks: Combining closure, bridging and proximities. *Regional Studies, 49*, 936–954.
Cohen, W., & Levinthal, D. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly, 35*, 128–152.
Cowan, R., & Jonard, N. (2003). The dynamics of collective invention. *Journal of Economic Behavior & Organization, 52*, 513–532.
Cowan, R., & Jonard, N. (2004). Network structure and the diffusion of knowledge. *Journal of Economic Dynamics & Control, 28*, 1557–1575.
Cowan, R., & Jonard, N. (2009). Knowledge portfolios and the organization of innovation networks. *Academy of Management Review, 34*, 320–334.
Cowan, R., Jonard, N., & Özman, M. (2004). Knowledge dynamics in a network industry. *Technological Forecasting and Social Change, 71*, 469–484.
Csaszar, F. A., & Siggelkow, N. (2010). How much to copy? Determinants of effective imitation breadth. *Organization Science, 21*, 661–676.
Fitjar, R. D., Huber, F., & Rodríguez-Pose, A. (2016). Not too close, not too far: Testing the Goldilocks principle of 'optimal' distance in innovation networks. *Industry and Innovation, 23*, 465–487.
Fleming, L., Chen, D., & Mingo, S. (2007). Collaborative brokerage, generative creativity, and creative success. *Administration Science Quarterly, 52*, 443–475.
Fleming, L., King, C., III, & Juda, A. I. (2007). Small worlds and regional innovation. *Organization Science, 18*, 938–954.
Gatti, C., Volpe, L., & Vagnani, G. (2015). Interdependence among productive activities: Implications for exploration and exploitation. *Journal of Business Research, 68*, 711–722.
Gilsing, V., Nooteboom, B., Van Haverbeke, W., Duysters, G., & Van den Oord, A. (2008). Network embeddedness and the exploration of novel technologies: Technological distance, betweenness centrality and density. *Research Policy, 37*, 1717–1731.
Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning.* Reading, MA: Addison-Wesley.

Granco, M., & Hoetker, G. (2009). NK Modeling methodology in the strategy literature: Bounded search on a rugged landscape. In D. D. Bergh, & D. J. Ketchen (Eds.), *Research methodology in strategy and management* (Vol. 5, pp. 237–268). Emerald: Group Publishing Limited.
Granovetter, M. (1985). Economic action and social structure: The problem of embeddedness. *American Journal of Sociology, 91*, 481–510.
Gupta, A., Smith, K., & Shalley, C. (2006). The interplay between exploration and exploitation. *Academy of Management Journal, 49*, 693–706.
Hamming, R. W. (1950). Error detecting and error correcting codes. *Bell System Technical Journal, 29*, 147–160.
Huber, H. (2012). On the role and interrelationship of spatial, social and cognitive proximity: Personal knowledge relationships of R&D workers in the Cambridge information technology cluster. *Regional Studies, 46*, 1169–1182.
Janssen, M. J., & Frenken, K. (2019). Cross-specialisation policy: Rationales and options for linking unrelated industries. *Cambridge Journal of Regions, Economy and Society, 12*, 195–212.
Kauffman, S. A. (1993). *The origins of order. Self-organization and selection in evolution*. New York and Oxford: Oxford University Press.
Laursen, K., Leone, M. I., & Torrisi, S. (2010). Technological exploration through licensing: New insights from the licensee's point of view. *Industrial and Corporate Change, 19*, 871–897.
Lavie, D., Stettner, U., & Tushman, M. L. (2010). Exploration and exploitation within and across organizations. *Academy of Management Annals, 4*, 109–155.
Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly, 52*, 667–694.
Levinthal, D. A. (1997). Adaptation on rugged landscapes. *Management Science, 31*, 934–950.
Lieberman, M. B., & Asaba, S. (2006). Why do firms imitate each other? *Academy of Management Review, 31*, 366–385.
Lissoni, F. (2001). Knowledge codification and the geography of innovation: The case of Brescia mechanical cluster. *Research Policy, 30*, 1479–1500.
Lissoni, F., Llerena, P., & Sanditov, B. (2013). Small worlds in networks of inventors and the role of academics: An analysis of France. *Industry and Innovation, 20*, 195–220.
March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science, 2*, 71–87.
Marsili, O. (2002). Technological regimes and sources of entrepreneurship. *Small Business Economics, 19*, 217–231.
Miller, R., Hobday, M., Leroux Demers, T., & Olleros, X. (1995). Innovation in complex systems: The care of flight simulators. *Industrial and Corporate Change, 4*, 363–400.
Miller, K. D., Zhao, M., & Calantone, R. J. (2006). Adding interpersonal learning and tacit knowledge to March's exploration-exploitation model. *The Academy of Management Journal, 49*, 709–722.
Muthukrishna, M., & Henrich, J. (2016). Innovation in the collective brain. *Philosophical Transactions of the Royal Society B: Biological Sciences, 371*, 20150192.
Nelson, R. R., & Winter, S. G. (1982). *An evolutionary theory of economic change*. Cambridge, MA: Harvard University Press.
Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Review, 45*, 167–256.
Nooteboom, B. (2000). Learning by interaction: Absorptive capacity, cognitive distance and governance. *Journal of Management and Governance, 4*, 69–92.
Nooteboom, B., Van Haverbeke, W., Duysters, G., Gilsing, V., & Van den Oord, A. (2007). Optimal cognitive distance and absorptive capacity. *Research Policy, 36*, 1016–1034.
Ordanini, A., Rubera, G., & De Fillippi, R. (2008). The many moods of inter-organizational imitation: A critical review. *International Journal of Management Reviews, 10*, 375–398.
Özman, M. (2009). Inter-firm networks and innovation: A survey of literature. *Economics of Innovation and New Technology, 18*, 39–67.
Phelps, C., Heidl, R., & Wadhwa, A. (2012). Knowledge, networks, and knowledge networks: A review and research agenda. *Journal of Management, 38*, 1115–1166.
Posen, H. E., Lee, J., & Yi, S. (2013). The power of imperfect imitation. *Strategic Management Journal, 34*, 149–164.
Powell, W. W., Koput, K. W., & Smith-Doerr, L. (1996). Interorganizational collaboration and the locus of innovation: Networks of learning in biotechnology. *Administrative Science Quarterly, 41*, 116–145.
Puranam, P., Stieglitz, N., Osman, M., & Pillutla, M. M. (2015). Modelling bounded rationality in organizations: Progress and prospects. *Academy of Management Annals, 9*, 337–392.
Pyka, A. (2002). Innovation networks in economics: From the incentive-based to the knowledge-based approaches. *European Journal of Innovation Management, 5*, 152–163.
Rivkin, J. W. (2000). Imitation of complex strategies. *Management Science, 46*, 824–844.
Savino, T., Messeni Petruzzelli, A., & Albino, V. (2017). Search and recombination process to innovate: A review of the empirical evidence and a research agenda. *International Journal of Management Reviews, 19*, 54–75.
Schilling, M. A., & Phelps, C. C. (2007). Interfirm collaboration networks: The impact of large-scale network structure on firm innovation. *Management Science, 53*, 1113–1126.
Simon, H. A. (1996). *The sciences of the artificial* (3rd ed.). Cambridge: The MIT Press.
Solís-Molina, M., Hernández-Espallardo, M., & Rodríguez-Orejuela, M. (2018). Performance implications of organizational ambidexterity versus specialization in exploitation or exploration: The role of absorptive capacity. *Journal of Business Research, 91*, 181–194.
Steinmo, M., & Rasmussen, E. (2016). How firms collaborate with public research organizations: The evolution of proximity dimensions in successful innovation projects. *Journal of Business Research, 69*, 1250–1259.
Stettner, U., & Lavie, D. (2014). Ambidexterity under scrutiny: Exploration and exploitation via internal organization, alliances, and acquisitions. *Strategic Management Journal, 35*, 1903–1929.
Uotila, J. (2018). Punctuated equilibrium or ambidexterity: Dynamics of incremental and radical organizational change over time. *Industrial and Corporate Change, 27*, 131–148.
Uzzi, B. (1996). The sources and consequences of embeddedness for the economic performance of organizations: The network effect. *American Sociological Review, 61*, 674–698.
Uzzi, B., Amaral, L. A. N., & Reed-Tsochas, F. (2007). Small-world networks and management science research: A review. *European Management Review, 4*, 77–91.
Van Rijnsoever, F. J., Van Den Berg, J., Koch, J., & Hekkert, M. P. (2015). Smart innovation policy: How network position and project composition affect the diversity of an emerging technology. *Research Policy, 44*, 1094–1107.
Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature, 393*, 440–442.
Wilden, R., Hohberger, J., Devinney, T. M., & Lavie, D. (2018). Revisiting James March (1991): Whither exploration and exploitation? *Strategic Organization, 16*, 352–369.
Wilensky, U. (1999). *NetLogo*. Available at: Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University http://ccl.northwestern.edu/netlogo/.
Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding, and selection in evolution. *Proceedings of the Sixth International Congress on Genetics, 1*, 356–366.
Wuyts, S., Colombo, M. G., Dutta, S., & Nooteboom, B. (2005). Empirical tests of optimal cognitive distance. *Journal of Economic Behavior and Organization, 58*, 277–302.