

A Complexity Perspective on Innovation in Ecosystems

Pieter J. den Hamer

© Pieter den Hamer, 2020
isbn: 978-90-830376-9-1

print: Print Service Ede

A Complexity Perspective on Innovation in Ecosystems

Een Complexiteitsperspectief op Innovatie in Ecosystemen
(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht op
gezag van de rector magnificus, prof. dr. H.R.B.M. Kummeling, ingevolge
het besluit van het college voor promoties in het openbaar te verdedigen
op vrijdag 18 december 2020 des middags te 4.15 uur

door

Pieter Jacobus den Hamer

geboren op 1 december 1967
te Barneveld

Promotor

Prof. dr. Koen Frenken

Table of content

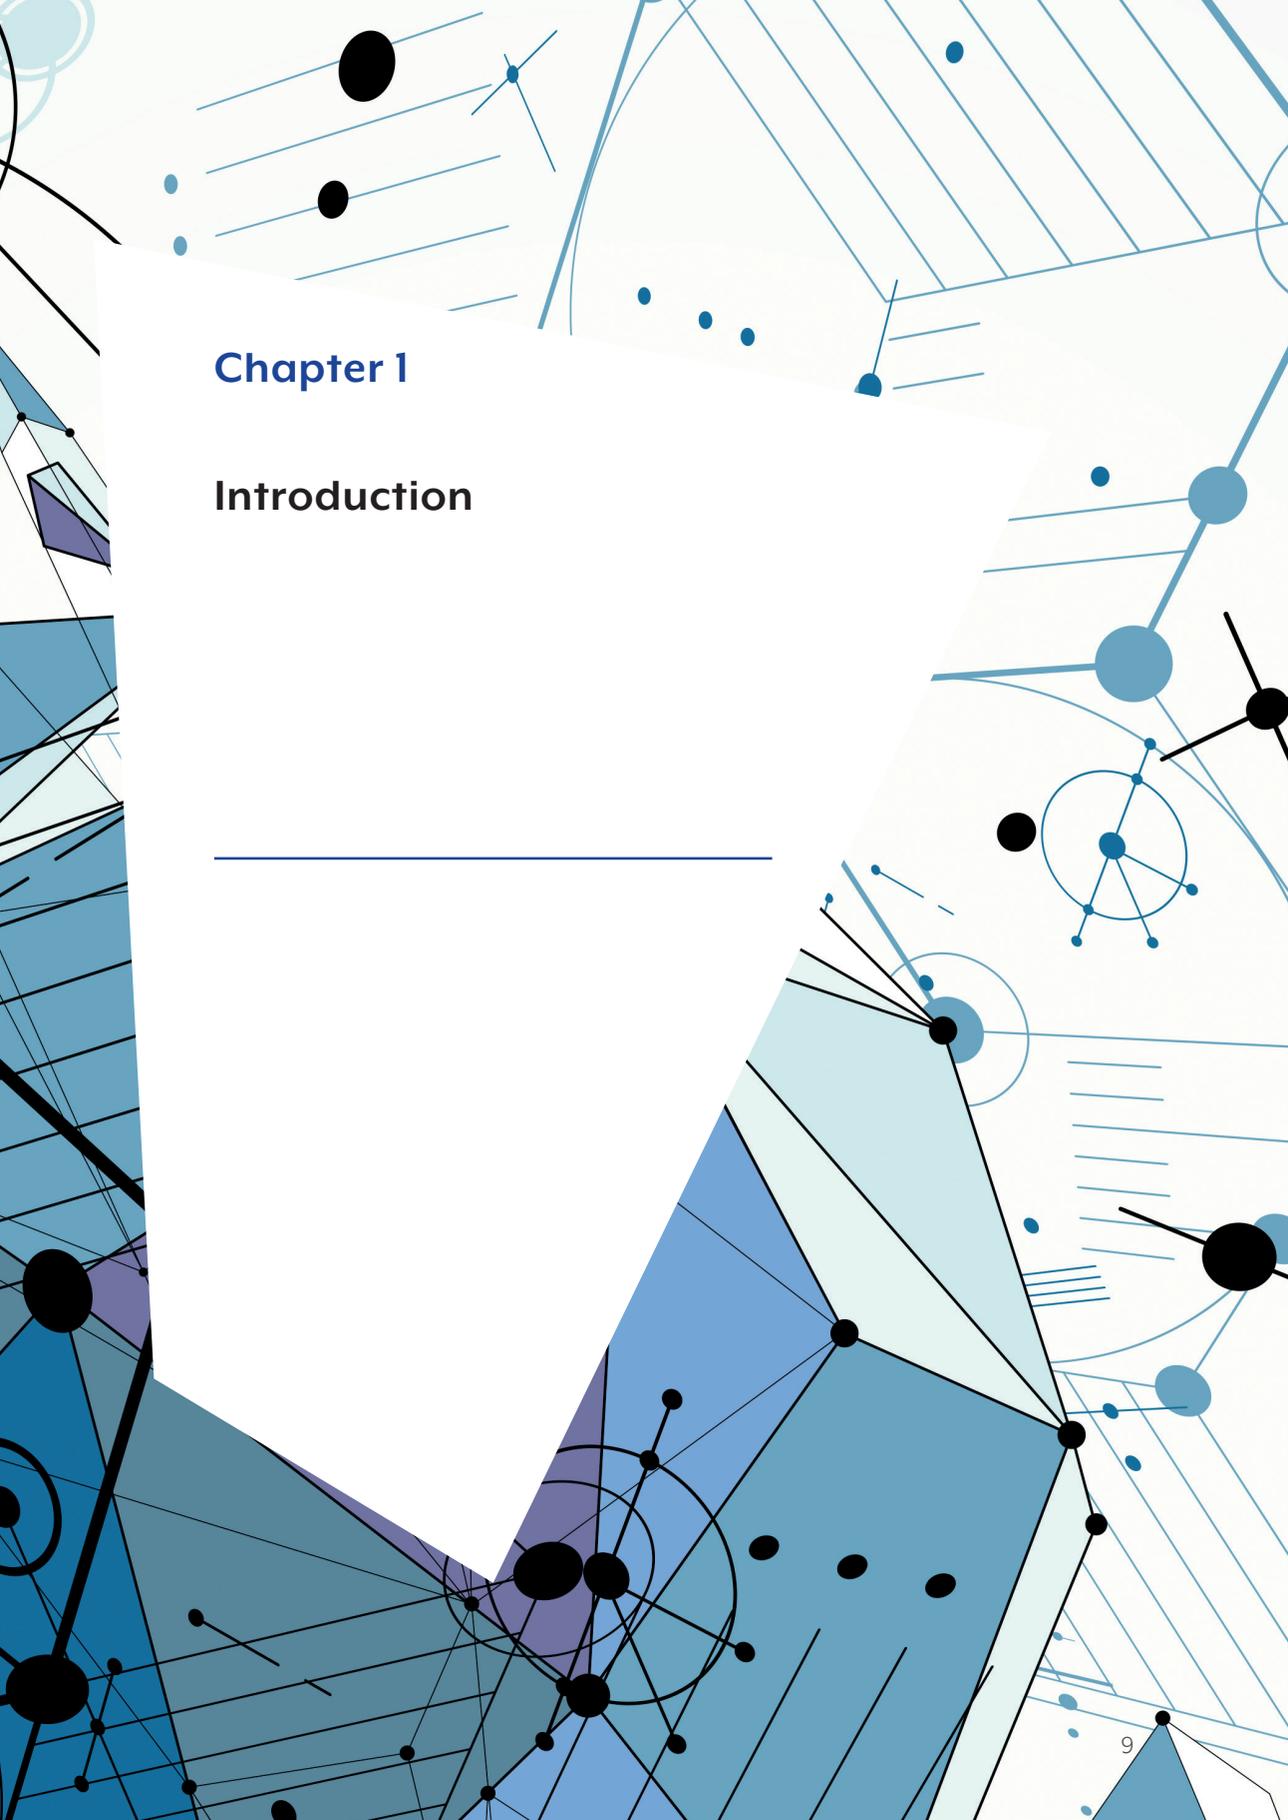
Chapter 1		
Introduction		9
1.1	Introduction	10
1.2	Innovation, complexity and ecosystems	11
1.3	An ecosystem perspective on innovation	13
1.4	Thesis outline	16
Chapter 2		
Balancing knowledge sharing in innovation contests		21
	Abstract	22
2.1	Introduction	23
2.2	Background	25
	2.2.1 Genetic algorithms	25
	2.2.2 Innovation challenge	30
	2.2.3 Innovation simulation	32
	2.2.4 Innovation modeling	33
	2.2.5 Simulations	34
2.3	Results	38
2.4	Conclusion	42
Chapter 3		
A network-based model of exploration and exploitation		45
	Abstract	46
3.1	Introduction	47
3.2	Theory	49
3.3	The Model	51
	3.3.1 Exploitation	52
	3.3.2 Exploration	54
	3.3.3 Simulation	57
3.4	Results	58
	3.4.1 Network-level	58
	3.4.2 Dyad-level	60
	3.4.3 Sensitivity analysis	62
3.5	Conclusion	63

Chapter 4		
Co-innovation in business ecosystems		67
	Abstract	68
4.1	Introduction	69
4.2	Background	70
4.3	The NKCS-model	73
4.4	Ecosystem regimes in NKCS	75
4.5	Model Implementation	77
4.6	Results	80
	4.6.1 Reproducing NKCS grids	82
	4.6.2 Network structures and ecosystem regimes	82
	4.6.3 Network structures and power laws	87
	4.6.4 Resilience	89
	4.6.5 Evolving ecosystem network structures	92
4.7	Discussion	94
	4.7.1 Implications	95
	4.7.2 Future work	97

Chapter 5		
A Complexity Perspective on Innovation Strategy in Business Ecosystems		99
	Abstract	100
5.1	Introduction	101
5.2	Innovation Strategy in Business Ecosystems	102
	5.2.1 Business Ecosystems	102
	5.2.2 Ecosystem Innovation Risks	104
	5.2.3 Innovation Strategy and Ecosystem Position	105
	5.2.4 Innovation Strategy & Ecosystem Reconfiguration	106
5.3	Complexity Perspective on Business Ecosystems	108
	5.3.1 Internal and External Complexity	108
	5.3.2 Ecosystem complexity model	111
5.4	Ecosystem positions	112
	5.4.1 Ecosystem positions and relations	116
5.5	Moving in the Ecosystem	119
	5.5.1 Complexity Strategies	119
	5.5.2 Mover Typology and Patterns	120
5.6	A Lifecycle Perspective on Ecosystem Positions	125
5.7	Conclusion	128

Chapter 6		
Conclusion		133
6.1	Introduction	134
6.2	Contest Innovation	135
	6.2.1 Summary of results	135
	6.2.2 Managerial & policy implications	135
	6.2.3 Limitations and future research	135
6.3	Market innovation	136
	6.3.1 Summary of results	136
	6.3.2 Managerial & policy implications	137
	6.3.3 Limitations and future research	137
6.4	Co-Innovation	139
	6.4.1 Summary of results	139
	6.4.2 Managerial & policy implications	139
	6.4.3 Limitations and future research	140
6.5	Combining perspectives: the ecosystem complexity model	140
	6.5.1 Summary of results	141
	6.5.2 Managerial & policy implications	142
	6.5.3 Limitations and future research	143
6.6	Overall conclusions	143
References		147
Summary		158
Samenvatting		162
Acknowledgements		166
Curriculum Vitae		168
List of publications (selection)		170





Chapter 1

Introduction

1.1 Introduction

Societies solve problems and seize opportunities, large and small, through innovation: the development and application of novel solutions, instantiated in new technologies products, services, capabilities, processes, business models or managerial, policy and governance practices (Schumpeter, 1939). Innovation is critical in all aspects of life. Countless innovations, from major breakthroughs to minor upgrades, all contribute to improvements in our welfare, wealth, health, security and safety, or to more profound transitions that are required to solve or mitigate large societal challenges like sustainability and climate change. Improving our understanding of how innovations come about in individual organizations (March, 1991), markets (Schumpeter, 1939) and ultimately societies (Freeman, 1987), is therefore highly relevant.

Most importantly, our further understanding of innovation must be accomplished in the context of an ever more complex world. Rising levels of complexity, resulting from global trends such as globalization and digitization, are reflected in more interdependencies between technologies, products, regions or firms, implying that innovation in one part or place is ever more likely to affect other parts or places (Dicken, 2015). And, *vice versa*, that innovation in one part is ever more frequently triggered by changes in other parts too. Increasingly, finding a new solution is not a matter of improving a single part or aspect, but requires the making of multiple, highly interdependent choices (Simon, 1962).

To understand innovation, we therefore must take the perspective of complexity – in terms of interdependencies between firms, products or components – into account (Simon, 1962; Simon, 1996), which is reflected in management science by a growing body of both theoretical and empirical research, with the work of Levinthal (1997) being an important milestone. Levinthal formalized the notion of innovation as an evolutionary ‘search’ process (March, 1991; Nelson & Winter, 1982), also metaphorically described as a ‘walk’ on a performance ‘landscape’ (Wright, 1932), using a complexity model. In turn, this formal model, known as the NK-model – describing ‘solutions’ as locations in a search landscape with N components or ‘choices’ and K interdependencies – was based on the seminal work in evolutionary biology by Kauffman (1993). Most notably, the model has proven to be an effective way of representing the difficulties in dealing with growing numbers of interdependencies, with higher complexity resulting in an ever more rugged landscape that becomes harder to navigate. In cases of highest complexity and most rugged landscapes, making a mistake in changing a carefully synthesized component assembly of good performance, is easily equivalent to falling from a high peak into a deep ravine (Levinthal, 1997; Simon, 1996; Wright, 1932).

Although other approaches to search or complex decision making with many interdependencies, also exist (Porter & Siggelkow, 2008; Puranam et al., 2015), the NK-model has become a generic model of search in management literature, systematically reviewed by Ganco and Hoetker (2008), Puranam et al. (2015) and Baumann et al. (2019). Building on the work of Kauffman (1993) and Levinthal (1997), scholars have focused on different aspects and variants of the search process, recognizing two main variants. In the first variant, known as ‘exploitation’ (March, 1991), innovations occur incrementally by changing components one at a time, gradually improving the

existing product or 'solution' (Baumann et al., 2019). This may occur by blind trial and error, or in a more knowledgeable way (Gavetti & Levinthal, 2000), for example by identifying improvements that are most beneficial to performance or other ways to limit the set of choices under consideration (Rivkin & Siggelkow, 2003). In the second variant, known as 'exploration' (March, 1991), innovators try to make a long 'jump' (Levinthal, 1997) in their landscape, simultaneously altering multiple parts of their design, either randomly or based on available insights (Gavetti & Levinthal, 2000), or by imitating another solution (Csaszar & Siggelkow, 2010; Lazer & Friedman, 2007; Posen & Martignoni, 2018; Rivkin, 2000). Other studies emphasize the effect of the architectural change of solutions (Ethiraj & Posen, 2014), especially the alteration of the structure of interdependencies between parts through modularization (Frenken & Mendritzki, 2012; Levinthal & Workiewicz, 2018). Still others pay attention to the time dynamics of the search process, such as the optimal order of search and the time needed to improve performance (Kauffman et al., 2000).

Notwithstanding the advancement of our theoretical understanding brought about by the application of the NK-model in the realm of innovation, it is striking that the large majority of models view innovation as a learning process taking place exclusively within the firm. This view stands in contrast to the burgeoning literature on R&D alliances (Oxley & Sampson, 2004), open innovation (Chesbrough, 2003a; 2003b), and business ecosystems (Adner, 2012). These literatures all stress, in different ways, that firms rely in many ways on other firms and public organizations in their innovation process. This begs the question how such inter-firm dependencies can be integrated in formal NK-models in such a way that not only internal and external learning processes are modelled jointly, but also the interrelations between internal and external organization and dynamics can be understood in a theoretical and analytical manner.

1.2 Innovation, complexity and ecosystems

From a complexity perspective, a logical starting point in addressing inter-organizational relations in innovation processes concerns the observation that interdependencies are not limited to the manifold choices made within the firms, but also exist between choices made in different organizations (Adner, 2012). In particular, technology choices made by firms in a supply chain often create challenges to other firms in the supply chain that they are part of. More generally, one can consider firms as being part of a business ecosystem, which can be defined as (Moore, 1993, p.76)

"An economic community supported by a foundation of interacting organizations and individuals – the organisms of the business world. The economic community produces goods and services of value to customers, who are themselves members of the ecosystem. The member organisms also include suppliers, lead producers, competitors, and other stakeholders."

Co-innovation by other actors that make up the ecosystem (Moore, 1993) in which organizations operate, must be taken into consideration in order to understand and manage innovation effectively (Adner & Kapoor, 2010; Adner, 2017; Iansiti & Levien, 2004). A growing number of case studies underline the importance of this ecosystem

perspective (Gawer & Cusumano, 2014a; Kapoor & Agarwal, 2017). For example, in the film industry (Adner, 2012), the introduction of digital film as new medium, required the co-innovation of a variety of organizations in the film ecosystem, such as studios and manufacturers of editing, camera and microphone tools and equipment. As another example, consider the ecosystem of the Windows operating system, with Microsoft as platform provider and a large number of complementor firms that offer products and services (Iansiti & Levien, 2004). Likewise, awareness of the innovation activities of competing ecosystem actors is important, not only to see if lessons can be learned from competitors (through imitation) (Rivkin, 2000) but also by co-developing technology or standards (through alliances) (Verspagen & Duysters, 2004).

However, despite the fact that there is a growing awareness about the relevance of business ecosystems (Moore, 1993), particularly in the context of innovation strategy (Adner, 2012; Iansiti & Levien, 2004), formal modeling on innovation still has focused mainly on innovation *within* organizations. Although a firm's own strategy, capabilities and internal operations are of course highly relevant to its innovation performance, it only seems obvious that the way a firm works together with its ecosystem partners, such as partners, clients and suppliers, and competes with other firms, is equally highly relevant. Innovation management and policies, then, can only be effective if we develop a further and deeper understanding of how firms depend on other firms and other ecosystem actors in their innovations, and how firms do not innovate in isolation but in collaboration or competition with other firms.

This thesis intends to contribute to this understanding by going beyond mere concepts and empirical case studies, which has been the emphasis of most extant work. To progress towards a deeper understanding of ecosystem innovation, this research takes the approach of modeling and simulations to reveal plausible mechanisms to explain a diversity of empirical findings and stylized facts. The current work seeks to contribute to management science by enriching our understanding of innovation by adding and applying the perspectives of ecosystems, complexity and evolution, with the aim to support and improve innovation management practices and policies. Considering this, it is clear that the current work relies heavily on insights from most notably complexity science and theoretical biology. As such, it fits in a broader trend that recognizes the need for applying theories and concepts that go beyond the more traditional reductionist view, both in science and practice. Although highly successful in many areas, reductionist sciences seem to overlook important phenomena, while considering systems as if in equilibrium state and by decomposing the system into its constituent parts, not taking dynamics and interdependencies into account (Beinhocker, 2007; Simon, 1996). Although useful, to gain further and deeper understanding of biological, human, economic, sociological or other complex systems, it is not sufficient to only decompose a system and analyze the workings of individual components. Instead, it is necessary to understand how components interact, in some cases leading to the emergence of higher order phenomena or regularities. Likewise, to understand the dynamics and time variance of complex systems at different scales, such as organizations, business ecosystems, markets and economies, it is increasingly clear that only focusing on the equilibrium states of such systems, does not provide the explanatory clarity that is required for insights about for instance the resilience or state transitions in such systems (Beinhocker, 2007; West, 2017). For this reason,

complexity science is increasingly being adopted as a cross-disciplinary paradigm in many scientific fields, including management science, economy and sociology, to which the current work aims to contribute.

1.3 An ecosystem perspective on innovation

To structure our thinking about innovation in an ecosystem context, we distinguish between two dimensions. First, it is considered that some innovations take place in a strongly coordinated fashion, often controlled and executed in the hierarchical context of a single organization or a tight consortium of firms, whereas other innovations take place in a looser coordinated fashion or solely within market context. Second, it is observed that some innovations are complementary to each other in the sense that these innovations may be combined to create synergy, whereas other innovations are more similar in their use and thus competing for the same resources.

	Single organization	Multiple organizations
Complementary	Conventional innovation	Co-innovation
Competing	Contest innovation	Market innovation

Figure 1-1. Ecosystem perspectives on innovation.

Combining these dimensions in a 2x2 matrix, results in four 'ecosystem perspectives on innovation' (see figure 1-1):

- **Conventional innovation.** More traditional or conventional ways of doing innovation are those that are coordinated, directed or orchestrated, often by a single organization. The work to realize the innovation is carried out according to some planning and agreed division of labor to exploit complementarities between different human skills and physical inputs (Levinthal, 1997). Such tight division of labor coupled with strict coordination may also be achieved with multiple firms as, for example, in joint ventures or R&D alliances (Verspagen & Duysters, 2004), coordinated by a consortium leader. Either way, the intended innovation is the result of collective efforts, with each unit contributing its own part of the whole that makes up the innovation.
- **Contest innovation.** An important disadvantage of the conventional way of doing internal innovation is that the search for a new solution is done in one way and in one way only. An alternative approach is to have multiple independent searches, increasing the likelihood that a more optimal solution can be found. To do so, an increasingly popular approach is to organize crowdsourced innovation or innovation contests (Adamczyk et al., 2012; Terwiesch & Xu, 2006), with an organizing firm inviting participants to enter into a competition with a financial or other incentive. Each participant is typically given the same innovation challenge. Participants may be internal to the organizing firm (competing teams of employees) or external to the organizing firm (as in crowd-sourcing). The latter

are often start-up firms that may see the contest as a tool for their market and business development.

- **Market innovation.** In a market economy, most innovation is fueled by competition and not explicitly organized by a single firm as in contest innovation, but rather in an unplanned manner with firms competing with each other for market share. Innovation on comparable new solutions will take place in competition between multiple firms in the market that are all seeking to improve their performance with similar products or services to meet existing or new customer needs. Being the first to launch a superior product or service helps firms to increase their market share at the expense of others (Nelson & Winter, 1982).
- **Co-innovation.** Finally, one can distinguish 'co-innovation' (Kauffman & MacReady, 1995; Kauffman et al., 2000). Increasingly, the development of new innovations becomes too large, risky or complex to be carried out by a single organization. Instead ecosystem actors may seek to collaborate with other actors, such as suppliers or clients, with each actor contributing by doing its own innovation on its own part (Adner, 2012). Together, all parts will make up a new solution, but without common ownership: all actors play their own independent role, often with the option to co-innovate with a choice of multiple complementary actors.

Of these perspectives, conventional innovation has been the main focus in the majority of past work. In contrast, the research that is presented in this thesis focuses on extending management science with insights about innovation in which the external ecosystem is most relevant, namely in the perspectives of contest innovation, market innovation and co-innovation.

A common theme in all three perspectives is if and how actors interact in their innovation activities. Such interactions may be coordinated or may take place with mutual consent, such as knowledge sharing and partnering, or uncoordinated or without mutual consent, such as imitation of competitor innovations or parts thereof. With or without coordination or consent, in all cases ecosystem actors could benefit from insights on which types and ways of interacting are most effective, either for individual firms or for the ecosystem as a whole. The first being the concern of innovation managers that work at company level, the latter being the concern of innovation policy makers in government or very large firms that dominate their ecosystem.

Following the scheme in Figure 1-1, and going beyond the analysis of conventional innovation, this thesis addresses three research questions, each corresponding with one of the three remaining perspectives:

1. For **contest innovation** (Adamczyk et al., 2012; Terwiesch & Xu, 2006): one approach for an organizing firm would be to have contestants not share any results or knowledge between them, with each participating firm following its own unique innovation path, resulting in multiple independent approaches to the same challenge. On the other hand, sharing intermediate results may prevent participants from spending too much time on an ineffective approach or running into a dead end, and may ultimately lead to a better winning approach. Therefore, under which considerations would it be useful for the contest organizing firm or

participants, to allow or promote knowledge sharing between participants, and to what extent?

2. For **market innovation** (Nelson & Winter, 1982): in an open market economy, most firms do their own innovation, while closely monitoring what their competitors are doing. In case a firm's competitor seems to achieve better results, a firm is likely to be tempted to see if and how the results of the superior competitor can be imitated (Rivkin, 2000). Imitation, however, can be hampered by a lack of a common knowledge base or limited capacity to absorb the competitor's expertise, or by possible incompatibilities between existing components of the firm's solution and new components that are being imitated from the competitor (Rivkin, 2000). Imitation from a competitor with common social and cognitive traits will be easier, on average, but may also provide insufficient 'newness' to improve performance significantly (Boschma, 2005). On the other hand, imitation from a 'strange' competitor may result in too many imitation errors, which especially in case of higher complexity levels may result in unacceptable performance degradation. Therefore, taking different levels of complexity into account, what is an optimal balance between newness and avoidance of error making?
3. For **co-innovation** (Kauffman & MacReady, 1995; Kauffman et al., 2000): a firm that has many external dependencies in its ecosystem may find itself constantly responding to new innovations by other ecosystem actors. Even to such an extent that no sustainable improvements are feasible, as ecosystem dynamics may be overwhelming and render any change obsolete before becoming effective. On the other hand, firms that have hardly any external dependencies but instead have taken all complexity internally, may find themselves inflexible and slow to respond to inevitable changes in their market and ecosystem. Therefore, which levels of complexity, internal and external, provide the right mix to balance agility with stability (Kauffman & Johnsen, 1991)? And what other factors are relevant to optimize a firm's position in its ecosystem for co-innovation? To what extent should the structure of the underlying ecosystem network be taken into account (Adner, 2017)? In addition, firms play complementary roles in co-innovation. For example, some are platform providers, upon which other more specialized firms may build specific solutions (Iansiti & Levien, 2004). Therefore, which role categories or position types can be identified in ecosystems and how can firms actively influence their own position, making their innovation strategy 'ecosystem aware'?

The relevance of the ecosystem seems to be most prominent in the co-innovation perspective, as firms and organizations rely heavily on each other in the supply and adoption of each firm's innovations. But also in market innovation, firms can benefit from other firms by imitating other firm's innovations, the effectiveness of which depending on factors such as social proximity in the network structure of their mutual ecosystem. And, in contest innovation, it seems useful to be at least aware of the fact that other ecosystem actors are involved in similar innovations, and with whom knowledge sharing is likely to be relevant.

So although one may argue that the relevance of the ecosystem characteristics varies with the perspective, in all cases the common key theme is that firms do not innovate in isolation and that it is vital to the effectiveness of innovation to be aware of a firm's ecosystem characteristics, such as the actors that are part of the ecosystem, their innovation activities, the structure in which they interact and depend on each other and the position of one's firm itself in one or more ecosystems (Adner, 2012; Adner, 2017; Iansiti & Levien, 2004). In fact, some firms even have an innovation strategy that does not take an ecosystem as a given, but which tries to influence or even shape the ecosystem to create optimal conditions for innovation.

1.4 Thesis outline

Extant work on innovation in ecosystems has mainly focused on conceptualization, demonstrated by empirical case studies. However, more profound theoretical insights to underpin concepts and empirical observations and derived stylized facts, are currently mostly lacking (Jacobides et al., 2018). To begin filling this gap and complement existing work, the research that is presented here focuses on simulation studies, developing and validating models that provide further insights about underlying mechanisms.

For all three perspectives in scope of the current work, simulation studies were carried out, mimicking innovation as an evolutionary process (Holland, 1992; Kauffman & Johnsen, 1991; Luo, 2018; Mäkinen & Dedehayir, 2014). In all studies, the subjects of innovation or 'candidate solutions' have been modelled as abstract strings of N components. Innovation, then, can be understood as a mutation in such a string changing the state of a component.

To simulate contest innovation, we choose to model the contest as a well-defined problem that has to be solved by competing teams. The well-defined nature of the problem in this context follows from the organized nature of the contest, where the firm organizing the contest knows what it looks for, and uses the contest to arrive at the solution in the shortest time possible.

To simulate market innovation, we conceive innovation as a more open ended process where firms aim to improve a product or service by small 'local' changes towards local optima. Here, we use the canonical NK model from biology (Kauffman, 1993), which is also widely used in management science (for a recent review, see Baumann et al., 2019). Competition in a market setting, then, stems from firms trying to serve the same customers and in their innovation process also, occasionally, try to imitate solutions of successful others (Nelson & Winter, 1982). What we specifically model is that the success of such imitations is not only harder for larger strings (parameter N), but also when such strings, representing complex products or services, have more interdependencies (parameter K) (Rivkin, 2000).

Finally, to simulate co-innovation, we also introduce external interdependencies (parameter C), being interdependencies between components of solutions from different ecosystem actors, building on the NKC-model of co-evolution introduced in the domain of biology (Kauffman & Johnsen, 1991). Figure 1-2 summarizes the three perspectives.

The outline of this thesis follows the scheme presented in Figure 1-2 with three simulation models presented in chapters 2 to 4. The first model is on contest innovation in chapter 2, the second perspective on market innovation in chapter 3, the third perspective on co-innovation in chapter 4. A more holistic and qualitative approach to all perspectives is presented in chapter 5, introducing a strategy framework for ecosystem innovation that addresses firm ecosystems positions, interdependencies and dynamics, amongst others, from a complexity and evolutionary perspective. Chapter 6 will finally present overall conclusions and reflections. The outline of the thesis can thus be represented as in Figure 1-3.

Following this scheme, the importance of taking complexity into account increases with each consecutive chapter. In the first perspective of contest innovation (chapter 2), the problem itself is not complex and the interaction structure between teams is taken as a given. Here, then, the question for the contest organizer holds – given the problem and interaction structure at hand – to what extent competing teams should disclose information during the contest as to speed up the finding of the solution.

	Single organization	Multiple organizations
Complementary	Conventional innovation	Co-innovation Model NKC
Competing	Contest innovation Model N	Market innovation Model NK

Figure 1-2. Simulation models used throughout the thesis.

Chapter 5	Single organization	Multiple organizations
Complementary	Conventional innovation	Co-innovation Chapter 4
Competing	Contest innovation Chapter 2	Market innovation Chapter 3

Figure 1-3. Ecosystem perspectives within scope of the current work, and corresponding chapters.

In the next chapter (chapter 3), we model innovation as complex problem-solving through local search on a fitness landscape (‘exploitation’). We add to this perspective the possibility for firms to imitate other firms allowing for longer jumps on the landscape (‘exploration’). These jumps are failure prone as imitation can go wrong, and the probability of such mistakes occurring is in turn modelled as depending on the distance between imitating and imitated firm in a (exogenous) social network.

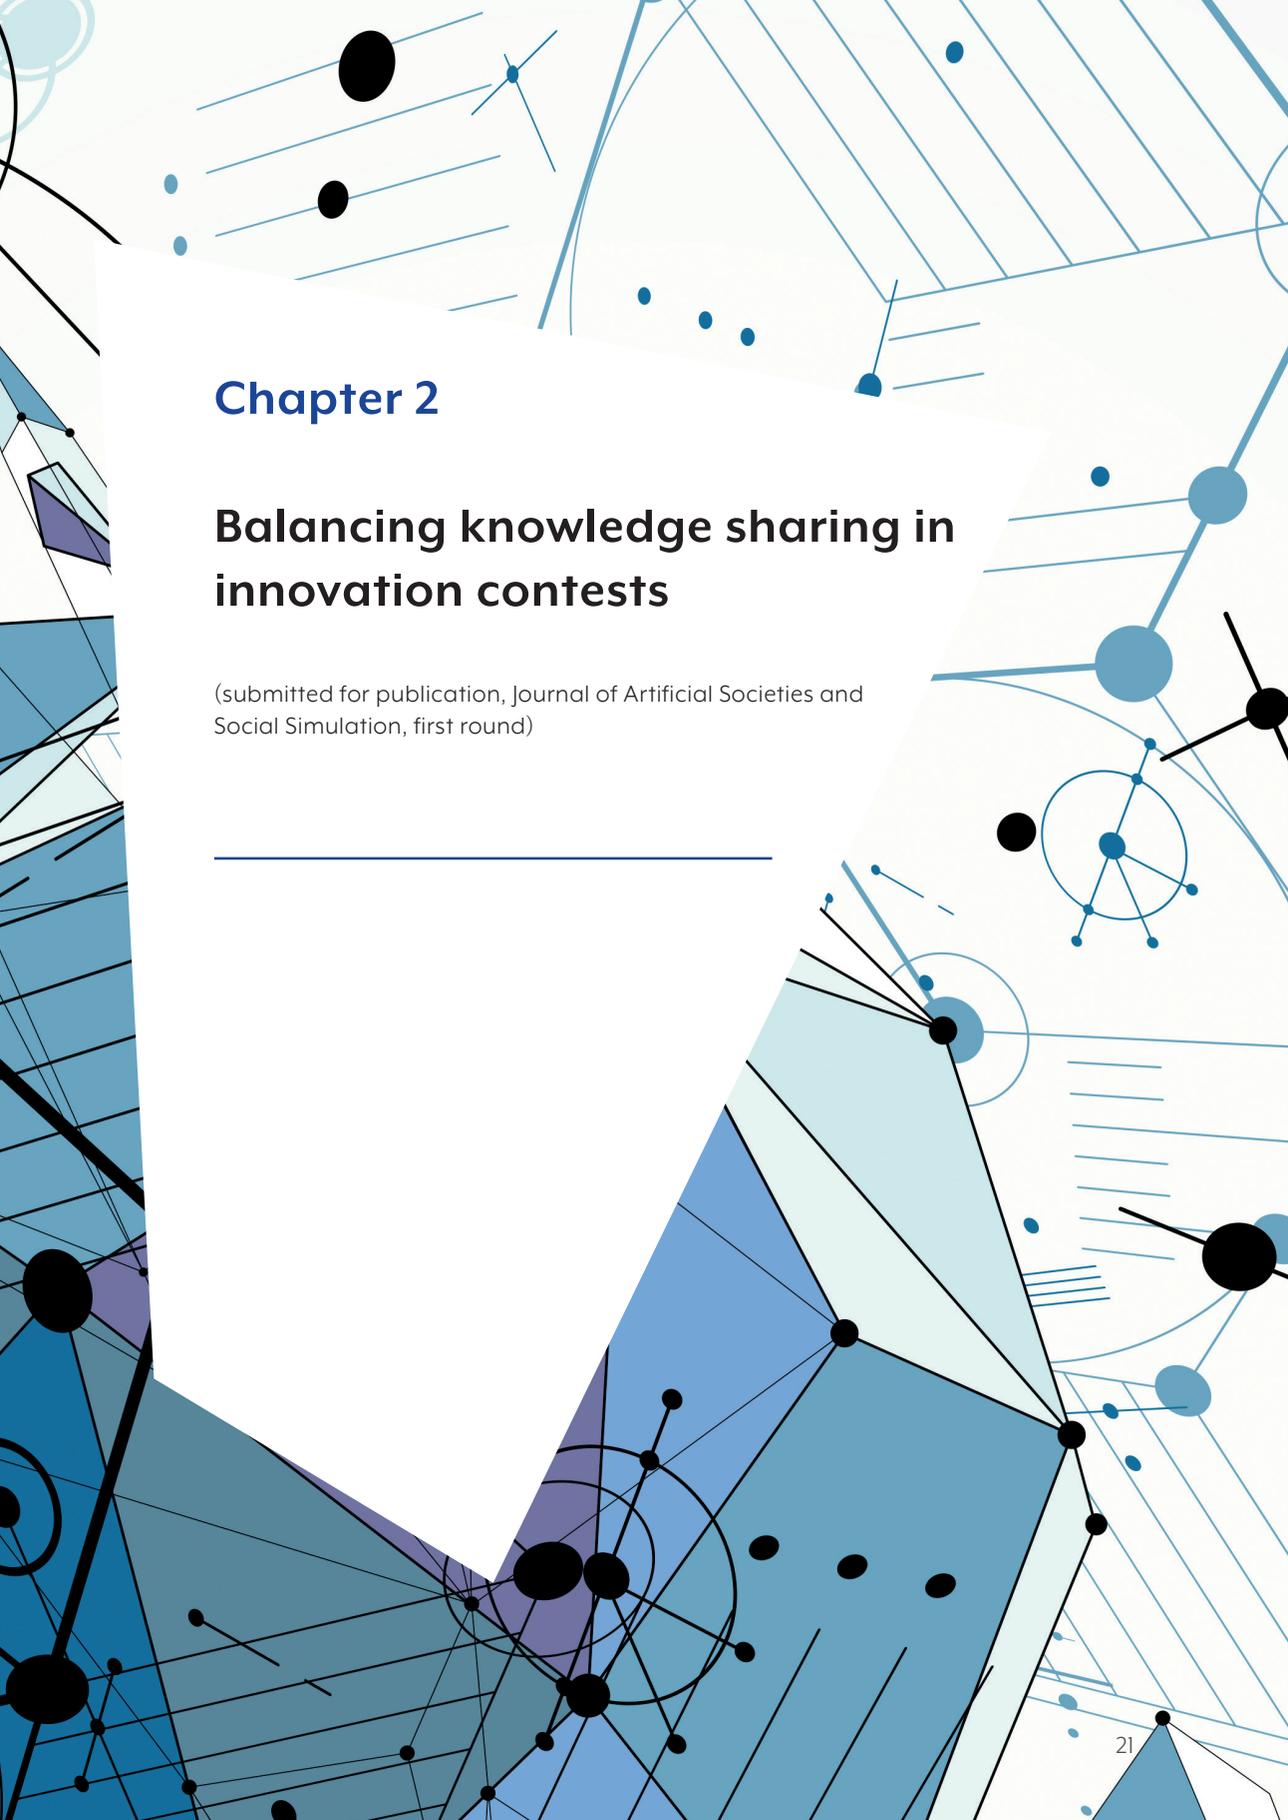
In the final modeling chapter (chapter 4), we turn to a fully-fledged ecosystem perspective where firms face interdependencies both internal to the firm and external to the firm. The ecosystem network itself defining the inter-firm dependencies is first considered as a given and, later on, also made endogenous.

The aim of each chapter on each perspective is to gain insights that should be relevant to both innovation managers and policy makers. Further to this end, the

perspectives are brought together in a strategy framework (chapter 5) to provide insights about firm positions and dynamics in ecosystems for innovation. All insights will be summarized in the final and concluding chapter, together with an overview of scientific results and suggestions for future work.







Chapter 2

Balancing knowledge sharing in innovation contests

(submitted for publication, Journal of Artificial Societies and Social Simulation, first round)

Abstract

Some firms organize innovation by having multiple teams pursue the same innovation goal in an innovation contest. This parallel approach to innovation may be more effective than a collective approach, in which a firm puts all resources into a single team, which runs the risk of dead ends. In innovation contests, however, the firm has to balance to extent to which competing firms share knowledge: little sharing tolerates teams to drift into unpromising areas, while excessive sharing may lead to premature convergence into a single solution, similar to the collective approach. In this simulation study, based on genetic algorithms, we analyze innovation costs while varying the rate and interval of knowledge sharing as well as different retention strategies that teams apply. The key and robust findings hold that the parallel approach outperforms the collective approach in almost all settings.

2.1 Introduction

When it comes to innovation, firms often choose to have their people and other resources working together in a common innovation endeavor (Tushman & Nadler, 1986), which we will call here the 'collective approach'. Obviously, this approach comes with the risk of innovation failure if a joint but single effort is too slow, takes a wrong turn or is unsuccessful altogether. An important alternative, therefore, is to have multiple teams compete to reach a certain innovation goal (Terwiesch & Xu, 2006), with each team having its own line of attack, which we will call here the 'parallel approach'. So, rather than putting all their eggs in one basket, firms may want to spread their bets to reach their innovation goals.

In a parallel approach, firms may staff multiple innovation teams with their own employees or outsource such task to external parties (possibly through crowdsourcing). Having multiple teams working in parallel on the same problem has been understood as an innovation contest (Adamczyk et al., 2012; Terwiesch & Xu, 2006). If outsourced, this is part of the trend towards 'open innovation' (Chesbrough, 2003a; 2003b; Terwiesch & Xu, 2006) in which external teams are invited to compete in solving a well-defined innovation problem, formulated by the firm that is organizing the innovation contest. This firm typically also provides incentives to attract teams for participation, which may include awards, working facilities, financing, branding, marketing, intellectual property sharing, co-licensing and other opportunities (Chesbrough, 2003a; 2003b). A well-known example is the DARPA Grand Challenge for autonomous robotic vehicles, in which teams from different companies compete to win a substantial amount of money, with criteria like speed, range and ability to deal with different terrain conditions (Seetharaman et al., 2006). Contests also occur in scientific research, where various scientific teams compete for priority, for example, to be the first to find proof for some theory (Franzoni & Sauermaun, 2014). Despite the ideal of scientific openness and transparency, often such teams keep their key knowledge to themselves or consciously delay sharing of their results, in order to secure their possible lead (Dasgupta & David, 1994).

Innovation contest may be organized by the initiating firm itself or by intermediary firms like for example InnoCentive, using specific or more generic crowd sourcing IT platforms (Haller et al., 2011), that include facilities to allow sharing of both intermediate and final results between participants and contest organizer, but also between participants. The latter becomes important when it is recognized that some knowledge or result sharing may be beneficial to reach a higher level of success (Gubbins et al., 2014), compared to the case where participants work strictly independent, being in direct competition and sharing as little knowledge as possible (Mitton, 2007).

A key question, in the case of the parallel approach, is how often and to what extent knowledge sharing between teams is optimal (Cowan et al., 2014; Gubbins et al., 2014; Nooteboom et al., 2007). On the one hand, too little knowledge sharing, possibly driven by the desire of contestants to protect their team's unique knowledge (Bogers, 2011) may result in lock-in effects (Nooteboom, 2000; Wuyts et al., 2005), meaning that a team's knowledge largely remains unchanged and is exploited to such an extent that no new innovations are possible any more. On the other hand,

too much knowledge sharing may result in loss of focus, as new knowledge that comes frequently or in large amounts, may distract teams from exploiting existing knowledge, preventing a team from converging towards useful innovation results, especially in more challenging contests (Hansen, 1999). Too much knowledge sharing (Weinberger et al., 2007) may also result in loss of originality, following the path of other teams prematurely, possibly leading to suboptimal performance or probably not very competing innovation results. Consequently, one must strike a balance between exploiting existing knowledge while preventing lock-in, versus exploring for new knowledge while preventing too much disruption of exploitation and premature follow-the-leader behavior. In finding this balance, relevant factors to optimize the sharing of knowledge should be actively identified and addressed, in the context of the parallel approach to innovation. Specifically: which knowledge, how often and in which quantities should be shared to reach the best innovation results?

What contrasts a parallel approach from a collective approach, is that in a collective search process all individuals work together, either by working in a single team or in multiple teams with a division of work, if possible. Either way, individuals are assumed to fully disclose their knowledge and share it with all other individuals at any moment in time. Of course, in practice, knowledge sharing, even within an organization, is challenging enough (Nonaka, 1991). Nevertheless, we assume that in the collective approach, people, capital, knowledge, equipment and other assets are shared, which saves on costs. Rather than spreading their resources thinly across parallel teams, a collective approach concentrates available resources, which may be the only feasible approach for small firms or for large-scale innovation projects. Then, the collective approach may be the only feasible way for firms that do not have sufficient resources for an internal parallel approach and do not want to work with external parties in an innovation contest.

However, the risk of the collective approach is that there is only a single dominant approach to solve the innovation challenge at hand, which may only be partially successful or may fail all together, without having an alternative approach readily available. In other words, because of the collective way of working and full sharing of knowledge, the collective may suffer from the risks of lock-in effects (March, 1991) or 'groupthink' (Postmes et al., 2001). In this context, groupthink is used as a concept to indicate the process in which people within the same social group may welcome and reinforce ideas that are more familiar, more trusted or more complementary to existing ideas, and may consciously or unconsciously inhibit new ideas that are very different from existing ideas and may require a paradigm shift (Postmes et al., 2001). Lock-in, in the sense of getting stuck in the innovation process on a suboptimal solution (March, 1991), may occur as a result of groupthink or for instance as a result of path-dependency, in which earlier investments or technological choices in the innovation process limit the search space of potential solutions. Exploring other areas would require significant redevelopments or, worst case, starting all over again.

Another downside of the collective approach may be the difficulty of initiating and coordinating such a collective way of working, with multiple stakeholders of for example different business units potentially having conflicting interests or overlapping or non-complementary assets. To collaborate, agreements are typically needed about the division of investments and potential but often unknown returns,

which in practice can be challenging. Likewise, planning and coordinating the work of many people and resources can be daunting – in general, large business activity programs and projects have a reputation of a high risk of failure (Martinsuo, 2013).

In contrast, in a parallel approach multiple teams work independently, each addressing the same innovation challenge in their own way, each using their own resources and assets. Sharing of knowledge is typically limited or may not take place at all. With this approach, the same ‘search space’ is explored in multiple directions, thus increasing the probability that at least one of the teams manages to find a proper solution, albeit with less available resources compared to the collective approach. That is, by comparing the results of the different teams, the most optimal solution can be selected by the jury of an innovation contest (Adamczyk et al., 2012).

A further advantage of the parallel approach holds that it requires much less or no coordination between teams and is quicker to initiate given the typically smaller size of the teams compared to the collective approach. Hence, organizing a contest may be a way to reduce the time to market. A well-known example of this kind of ‘coordinated parallel approach’ is the case of Siemens in which multiple independent teams within the same organization competed to imitate laser technology from the U.S. and were successful in doing so in only nine months (Albrecht, 1997; Schepers et al., 1999).

A key question in case of the parallel approach, then, is to find an optimal balance regarding the knowledge that teams should share during the search process. In one extreme, teams may never share knowledge during the innovation process, which would make sense from a competitive perspective and from the perspective that teams should not be influenced by other ideas and should stay committed to explore their specific area of the problem search space. However, in this extreme case, the implication could be that individual teams find themselves locked-in in an area that only contains suboptimal solutions, lacking the knowledge that others could provide to them to escape and find better solutions. In case of the other extreme, in which teams share all knowledge all the time, the parallel approach becomes equivalent to the collective approach, removing the benefits of exploring the solution space in multiple areas.

To investigate the optimal balance of knowledge sharing between teams involved in an innovation contest, we developed a model based on a genetic algorithm. The model specifies the amount and frequency of knowledge sharing between firms, and the strategies they apply regarding knowledge retention. The next section introduces this model, while section 3 discuss the results. We end with some conclusions and reflections as well as a discussion of the limitations of our study.

2.2 Background

2.2.1 Genetic algorithms

We model innovation as an evolutionary process (Fleming & Sorenson, 2001b), in which each team starts with a limited set of candidate solutions, which over time are improved by gradual changes and by interchanging parts of solutions that together result in new and better solutions, replacing inferior solutions. Typically, this process

starts out with a diverse set of candidate solutions, which will converge to a more limited set of solutions and will finally result in a single dominant solution, for each innovating team.

More specifically, we choose to model innovation by doing simulations with the application of genetic algorithms. These are computational models based on principles derived from population genetics in nature and inspired by Darwin's theory of evolution (1859). John Holland (1975; 1992) was amongst the first to study genetic algorithms. After his seminal work, genetic algorithms as general problem solvers have been applied to a wide variety of problems; both in science – for example to test biological theories on evolution and genetics – and in industry – for example to optimize logistics planning or production scheduling. What is more, genetic algorithms have been used conceptually as a model of how firms innovate and learn from each other (Birchenhall et al., 1997).

A standard genetic algorithm (Goldberg, 1989) starts out with a set of randomly generated, possible solutions. This set is called a population and has a fixed size. Possible solutions are called individuals, which are encoded as binary strings of equal length. According to the goodness of the solution, a fitness value is assigned to each individual. In proportion to the fitness value, individuals are then selected to reproduce themselves. Most 'offspring' is created by merging two parent individuals, using a cross-over operator. Also a genetic operator called mutation may be applied to a small part of the newly created population. Mutation implies that one or more bits in the binary string are flipped (from 0 to 1 or from 1 to 0). After mutation, the existing population is completely replaced by the new population. This process of evaluation (assigning fitness values), selection (according to fitness values), reproduction (cross-over) and variation (mutation) is iterated for a number of time steps, until some termination criterion has been met. Since individuals with a higher fitness value are more likely to be selected to reproduce, the genetic algorithm will typically converge after a number of iterations or generations to a single solution. This solution might be optimal or suboptimal, the latter meaning that the genetic algorithm has (prematurely) converged to a solution which locally in the search space might be best, but worse than other solutions that might be available if we could search the search space of solutions globally.

The standard genetic algorithm has a single population of possible solutions. Alternatively, the population is distributed over a number of subpopulations or 'islands', with each island carrying out its own local standard genetic algorithm, and with islands interchanging or 'migrating' a certain number of individuals, selected according to their fitness or other criteria, with a certain frequency. Pettey, Leuze and Grefenstette (1987) were among the first to publish about this kind of 'parallel genetic algorithm'.

Important characteristics that influence the performance of parallel genetic algorithms are the topology of the 'archipelago' or more precisely, the availability of connections between islands, population size (in total and for each island), migration interval, migration rate (Goldberg et al., 1995) and migration type, with the latter often ignored or implicit in previous work. When comparing different model implementations (Cantu-Paz, 1995) and their performance, it is clear that in general, especially for harder problems, parallel genetic algorithms are quite robust

in outperforming the standard or serial genetic algorithm. In other words, it seems to pay off to explore different areas of the search space, if large and rugged enough, and to prevent too early convergence on a specific area, even if parameters are not fine-tuned for optimal performance. It should be noted though that island populations have a critical minimum size, below which performance is very poor (Goldberg et al., 1995).

Migration is the process of moving individuals from one island to another, which is controlled by topology, migration interval, rate and type. The topology reflects which islands are connected and therefore between which islands migration may take place. Without loss of generality, most parallel genetic algorithm implementations assume that migration only takes place between islands that are directly connected. The migration interval controls how often or after how many time-steps or generations, migration takes place. The migration rate controls the number of individuals that is selected to be emigrated each time migration occurs. The type of migration indicates which individuals in the island population are selected to be emigrants, whether emigrants are cloned or relocated, and which individuals are selected to be replaced by immigrants. Replacement is necessary when – like in most parallel genetic algorithm implementations – the population size on each island must be kept constant and is not allowed to grow or shrink; both on islands of origin and on islands of destination.

When considering migration characteristics or parameters, it is not directly obvious how these may interact. One way of trying to understand their interactions is to look at the measure of isolation and the notions of local convergence versus global convergence. Here local convergence is defined as the convergence of the population of a single island towards a single local solution. Global convergence is defined as the convergence of all islands towards a single global solution.

Highly isolated islands, not so much disturbed by fresh genetic material from neighbors, will tend to locally converge towards their own dominant solution. Each island gets the opportunity to investigate its own part of the search space¹, resulting in a quite diverse exploration of the search space by all islands together, increasing the likelihood that at global level, after exchanging solutions between the different islands, a higher quality solution will be found. In case of less isolated islands, global convergence occurs more quickly. Yet a quick global convergence implies a quick local convergence, which may prevent islands from a more thorough exploration of their part of the search space (i.e. premature local convergence).

To put it differently, we may compare the case of highly isolated islands to using a shotgun: many but small bullets are likely to hit a target, provided that these bullets are still big enough to have sufficient impact. The other extreme case of having no isolation at all – equivalent to one large population on a ‘continent’ – can be compared to using a single-bullet rifle, which has a larger impact but only in one area of the target. It should be noted that this metaphor only holds after convergence has started to set in, because

¹ However, a smaller island population size (compared to a larger population size in case of a serial genetic algorithm or in case of fewer islands) may lead to a quicker local convergence, having a negative impact on the quality of found solutions on each island, on average. Therefore, in distributing a total population of a certain size over a number of islands, a balance must be found between the number of islands (more = better, but smaller population size) and the population size of each island (bigger = better, but less islands). This challenge is outside our current scope, but was addressed by carrying out preliminary experiments before simulations took place.

at the beginning of either case, the population is crowded with a diversity solutions, spread over the search space. But depending on the speed of convergence, this diversity will diminish and limit search to either a number of areas that is equal to the number of islands, or a single location when there is only a continent.

Table 2-1 summarizes how isolation is qualitatively influenced by different migration characteristics. When considering topology, it is clear that isolation is higher when there are fewer connections between islands. Isolation is also higher when islands are more clustered, i.e. when islands that are connected to one island are also more often connected to each other, resulting in slower migration on average in the island network as a whole, with longer average shortest path distances between islands in more clustered networks. In contrast, the average shortest path decreases in more random networks (Watts & Strogatz, 1998), thus reducing the number of intermediate islands to migrate from one island to another, with individual migration steps only occurring between directly connected islands.

Migration parameter	Isolation high	Isolation low
Topology	SParsely connected Highly clustered	Densely connected Random connections
Migration rate	Low	High
Migration interval	High	Low
Migration type	Move random individuals	Copy best individuals

Table 2-1. Qualitative effect of migration parameters on isolation

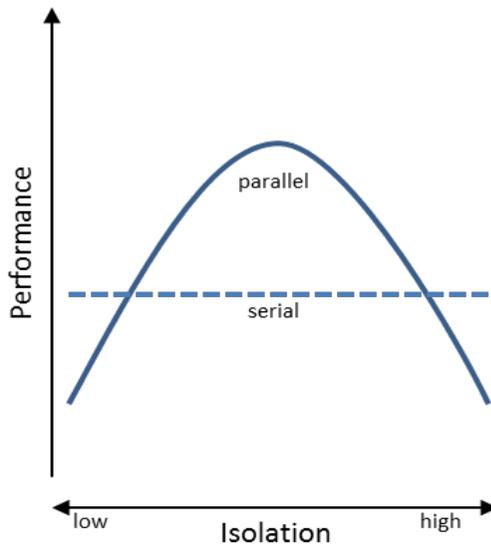


Figure 2-1. Expected effect of isolation on performance: very high isolation is equivalent to many small serial approaches, very low isolation is almost equivalent to a serial approach. In both extremes the parallel approach may become inferior to the serial approach. In between, the parallel approach is expected to outperform the serial approach.

Obviously, isolation also increases with lower volume or less frequent migration. Less obvious is the impact of the migration type on isolation: when individuals are randomly selected for migration for each migration step, it will take longer before individuals may reach other islands that are indirectly connected to their island of origin, requiring multiple steps in which the individual was selected at random. Selection of the same individual is more likely when selection is based on performance – the better the performance, the higher the probability of selection. In addition, when individuals are moved rather than copied, it will take longer before individuals and their offspring, created with cross-over, may increase their presence on multiple islands, creating the effect of more isolated islands.

The extent of isolation – together with the population size per island – determines the speed of local and global convergence and the probability of success in finding a highly performing solution. In general, slow convergence results in a higher probability of success, provided that the problem is difficult enough to require a more thorough search in a more rugged landscape, and *vice versa*. To control the speed of convergence and the probability of success to reach a certain performance level, both the size of the population and the extent of isolation are important factors. Other factors include the way individuals are selected to reproduce and the rate of mutation that is applied. But when comparing serial and parallel genetic algorithms, population size and isolation are the most important differentiators.

Both in serial and parallel genetic algorithms, a larger population size leads to a slower speed of convergence and a higher probability to reach higher performance levels, and *vice versa*. The effect of isolation, which is only available in parallel genetic algorithms, is more subtle. Figure 2-1 shows the qualitative effect of the measure of isolation on performance, for both parallel and standard or ‘serial’ genetic algorithm. The latter is insensitive to isolation, as migration doesn’t play a role in that case. Between the two extremes of very low and very high isolation is an area where parallel genetic algorithms are expected to reach higher performance levels than serial genetic algorithms. The main reason being that parallel genetic algorithms allow exploration of multiple areas of the search space for a longer period of time, in essence by using isolation to delay convergence, but still with the same total population size as the serial genetic algorithm. Each island does its own search, while migration provides result sharing and escapes from lock-ins.

However, performance of parallel genetic algorithms with very high or very low isolation tend to become lower. Implementations with very low isolation are basically similar to the serial case: when migration is very intense between islands, the different island populations become almost equivalent to a single integrated population. Performance will therefore be equivalent, or worse when local convergence at the level of individual islands leads to quicker global convergence. In the other extreme, where isolation is high, the processes on individual islands are equivalent to an equal number of serial but smaller sized approaches. However, because of smaller population sizes, providing less possibilities to explore the search space, average performance will tend to be lower, compared to the case where all island populations would have been combined in a single population in a serial approach. Or, in other words, when isolation is very high, islands will be locked-in: stuck on their local optimum, without having an escape.

Consequently, the parallel approach is not a panacea for all situations and will only be beneficial if the right balance is found between the characteristics that control migration and the resulting extent of isolation. As mentioned before, other work (Cantu-Paz, 1995) indicates that parallel genetic algorithms seem to be quite robust in terms of outperforming serial genetic algorithms, provided that solutions are not too trivial to render differences inconsequential.

2.2.2 Innovation challenge

To determine how different parameter settings of our model influence performance, we have run a series of innovation contest simulations. The question, then, holds how to mathematically represent the search space of an innovation contest. Following the seminal work by Levinthal (1997), it has become customary to represent the search space of an innovating agent as a complex N-dimensional fitness landscape, where the agent has to choose between two options (0 or 1) for each dimension making up a search space of size 2^N . The fitness of each component, however, is not just dependent on its own state, but also on the state of K other components yielding a 'rugged' fitness landscape with many local optima (for an extensive review of such NK-models in management studies, see Baumann et al., 2019).

To model an innovation contest, however, one may resort to simpler optimization problems as the problem to be solved is well-defined. In fact, different from innovation processes that venture into the unknown, a contest is an innovation process where the desired outcome is defined ex ante. Here, we chose the problem of 'counting linear weighted ones' as representing the innovation contest. The relative simplicity of this challenge reflects the fact that in innovation contests, the problem that needs to be solved by doing innovation, is well formulated, with the contest organizer knowing how to identify a successful solution. Solutions are often a matter of finding or developing the right components within an existing architecture, with a clear understanding if a component is meeting success criteria, or not. Think here for example of an organization that wants to imitate a product of a competitor, like in the case of Siemens who wanted to imitate a laser product of a competitor (Albrecht, 1997) or the more contemporary example of a platform organization like Apple that selects apps among supplier firms (Gawer & Cusumano, 2014a; 2014b).

In addition, the problem of 'counting linear weighted ones' reflects the fact that in such innovation contests, some components are more important than other. Consider, for example, an innovation contest to improve a small single-engine aircraft - with the firm that organizes the contest wanting to improve current aircraft models, either their own or those of competitors. It therefore requires that the aircraft should be designed according to the dominant architecture that defines the topology of tail, fuselage, wings, landing gear, engine, cockpit, interior, et cetera. The challenge then is to find for each of these components a better alternative. The dependencies between the parts that make up an aircraft is a given, in this context. The contestants are not asked to invent a new type of aircraft, but to improve or find the best possible parts for an existing aircraft type. However, some parts are more critical to its performance than others. For example, wings and engine are obviously more critical than seats and interior upholstery.

The problem of ‘counting linear weighted ones’ is about creating a string of bits (0 or 1) with as many 1s as possible, with the importance of bits being dependent on their position in the string. The further right or to the end of the string, the more weight each bit carries. For example, string ‘100111’ is better than ‘000111’ is better than ‘111000’. The actual ‘performance’ P_{actual} of a bitstring of length n (containing bits b_1, \dots, b_n) is calculated accordingly:

$$P_{actual} = \sum_{i=1}^n b_i \cdot i$$

with maximum performance being

$$P_{max} = \binom{n}{2} \cdot (n + 1)$$

The difficulty of the problem in the case of genetic algorithms is mainly dependent on the length of the bitstring. Preliminary experiments demonstrated that the difficulty of the problem is a relevant factor to consider. The easier the problem, the less differences exist between serial and parallel and between different parameter settings. To demonstrate this, we conducted our simulations with bitstrings of lengths 20, 50, 100, 200 and 500. Respectively with $P_{max} = 210$, $P_{max} = 1275$, $P_{max} = 5050$, $P_{max} = 20100$ and $P_{max} = 125250$. To compare results, we calculated the (relative) performance P as follows:

$$P = \frac{P_{actual}}{P_{max}}$$

In the context of innovation, we may consider a bitstring as a representation of a list of possible components or modules present (1) or absent (0) in a product or service, with components or modules ordered in order of importance (in terms of effect on performance), starting with least important at the top and ending with most important at the bottom of the list. For example, bitstring ‘1010110101’ may represent a list of an extremely simple airplane containing 10 possible components, with the first 1 representing the presence of a company emblem on the tail, with negligible effect on the performance level of the airplane, and the last 1 representing the presence of wings, which tend to be a defining characteristic of airplanes.

Finding as many 1s as possible is therefore equivalent to finding as many as possible performance improving modules that make up the product, service or process, assuming that no interdependencies between components exist. In this way, we can abstract from the possibility that flipping a zero into a 1 may decrease overall performance due to interdependencies between components as in the NK-model (Kauffman, 1993). Instead, we focus solely on the weighted sum of 1s as the single well-defined performance criterion that defines the context that competing teams are participate in. The goal, then, is to develop as many right-hand side components, in order of importance and as quick as possible.

In addition, an interesting aspect of the counting linear ones problem is the fact that components that at least initially are less important (more to the left in a bitstring) may 'disappear', especially in case of limited capacity in the contesting firms. By initially prioritizing solutions that have the most of the most important components (more 1s on the right) and that may have less of the less important components, the latter may be lost to the firm at hand. This, as simulations will demonstrate, can be remedied by exchanging knowledge between firms, further underpinning why and how knowledge sharing can be beneficial.

Despite the benefits of working with the counting linear ones problem, we must of course be careful of not assuming too easily that our findings will be applicable in real-world innovation settings. Still, our modeling and simulation findings do provide guidance to future empirical work to focus on relevant aspects in the sharing of knowledge between innovation contestants.

2.2.3 Innovation simulation

Simulation of the innovation process takes place by executing genetic algorithms as follows. Initially, a set of possible solutions or bitstrings of equal length is randomly initiated, with each bitstring having a maximum of 10% of 1s at random positions. Increasing this percentage would make the innovation challenge easier, and vice versa. Then, by applying cross-over operators, a new set of bitstrings is created – a process that is repeated for a number of time steps until convergence has been accomplished or a predefined maximum number of time steps has been reached.

Crossover is an operator to produce new solutions or 'children' by recombining existing solutions or 'parents' (Goldberg, 1989). In its most basic form, crossover is used with a single crossover point, applied to two parents, to select the part of parent 1 to the left of the crossover point that goes to child 2 and the part of parent 1 to the right of the crossover point that goes to child 1, and vice versa. See Figure 2-2 for an example.

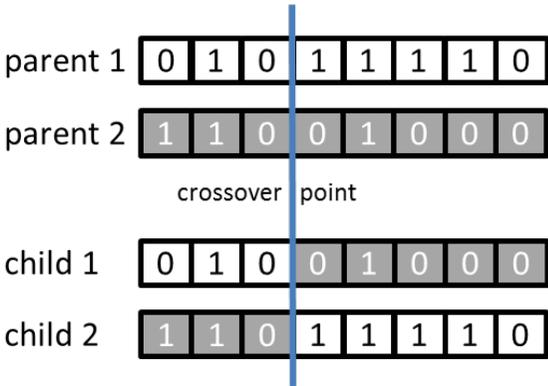


Figure 2-2. Example of crossover with a single crossover point.

To implement crossover, the elitist recombination (Thierens & Goldberg, 1994) technique was used, which is a variant to crossover in the standard genetic algorithm, integrating selection and recombination. In elitist recombination, two 'parent' solutions are randomly selected –without replacement– to produce two new solutions using crossover. Of the resulting four solutions (2 'parents' + 2 'children'), the two best performing solutions are then selected, according to their performance, to be inserted in the set of solutions for the next time step. In the example of Figure 2-3 and for the 'counting linear weighted ones' problem, parent 1 and child 2 offer best performance and are therefore selected to be inserted in the next solution set. The process of selecting 'parents', producing 'children' and then inserting the best performing solutions, is continued until the solution set for the next time step is completely filled. Since selected parents are removed from the existing solution set, parents only have a single opportunity to reproduce. At least the parent itself or one of its offspring is inserted into the solution set of the next time step. In pseudo code elitist recombination works as follows:

```

loop until the next solution set is filled
  randomly select and remove two parents from the current solution set
  produce two children using crossover
  insert best two of parents and children into the next solution set
end loop

```

The probability of crossover was 1: crossover was always applied. Crossover points were randomly selected. Mutation nor other genetic operators were used.

2.2.4 Innovation modeling

In many cases, the objective of applying parallel genetic algorithms has been to shorten the time duration of the algorithm to complete, by making use of parallel computing hardware. In other cases parallel genetic algorithms were inspired by or used to simulate the biological evolutionary concepts of punctuated equilibria based on allopatric speciation (diverging evolution of species after being geographically separated) and stasis (persistence of the genetic composition of a species after having reached equilibrium in a stable environment) (Cohoon et al., 1987).

In our case the objective is to use parallel genetic algorithms for modeling a parallel approach to innovation. For this, we model teams as 'islands', sharing their innovation results with other teams through connections with other islands. Each team has its own set of possible solutions for the innovation challenge at hand, that are being explored and improved. Poor performing solutions are discarded, better ones are kept, further improved upon, or combined with other solutions. Now and then, knowledge about one or more solutions is shared with other teams, introducing 'fresh blood' into the solution set of each team, possibly providing a new approach to improve or even make a step-change in solution performance – the newly injected solution in a sense providing an escape from a local optimum in which a team knowingly or unknowingly might have gotten stuck.

In this context, our main question is how innovation performance is affected by migration parameters such as the rate and interval of knowledge sharing, and which

knowledge is selected to be shared. Regarding the latter, we discern three cases: moving knowledge about random solutions or 'move random'; moving knowledge about the best solutions or 'move best'; and copying knowledge about the best solutions or 'copy best'. The first case of 'move random' reflects a situation in which knowledge gets shared without coordination, for example by employees that quit their job and move to another team or employer, taking their knowledge with them. The second case of 'move best' is equivalent to for example mandatory job rotation within a larger organization, requiring employees to change their role and/or team periodically, in order to create more of a learning experience or to prevent 'group think' or other negative effects of the same group of people working too long together (Postmes et al., 2001). We assume here that the knowledge that is being moved is 'tacit', meaning that it is hard to make explicit and share with others – when people move, their tacit knowledge goes with them. In the third case of 'copy best', on the other hand, knowledge can be codified and therefore more easily copied and shared with others, without losing it. This means that in the third case, knowledge stays on the island after it has been copied and shared with another island, and is not replaced by other knowledge. In contrast, in cases one and two, knowledge that is shared is not copied and therefore replaced by new knowledge from other teams that share their knowledge in their turn.

By doing simulations of innovation processes with parallel genetic algorithms and varying migration parameters, we are able to measure performance and time effects. Moreover, our primary interest goes to migration rate, interval and type, their interactions and their relative weight in terms of impact on performance and time. For the sake of clarity, we keep the topology of how organizations or units are connected, straightforward and constant. Varying network topology and understanding how topology influences knowledge sharing and innovation results, lies outside our current scope and is addressed in other work.

Performance is taken here to represent the overall quality of the subject of innovation, which may be related to technical quality, customer perceived quality, production costs or other quality factors and their combinations, to which our study is agnostic. In addition to performance we also consider time or the duration of the innovation process as an important indicator about the innovation process. From a business perspective, time is often related to costs – the longer innovation takes, the more resources & materials are needed, in general. But perhaps more importantly, the duration of the innovation process is also highly relevant when it comes to be amongst the first to enter an existing or new market with new products or services resulting from innovation. Being a 'first mover' or 'early mover' in terms of bringing a new product or service to the market, offers important advantages, including market share attainment and price determination, although it should be noted that first mover advantages may be reduced by other possibly negative effects that result from path dependency and lock-in (Querbes & Frenken, 2012). Nevertheless, most commercial organizations consider their 'time to market' as an important factor in their competitive strength and market position.

2.2.5 Simulations

To run our simulations, we determined by preliminary experiments that a total solution set size of 100 would be sufficient for our current goals, in the parallel case distributed over 10 teams, each with a solution set size of 10. Table 2-2 provides an

Parameter	Description	Setting(s)
P	Number of possible solutions = size of solution set or 'population size'	100 in total, 10 per team in parallel version
T	Number of teams	10 in parallel version, 1 in serial version
N	Length of bitstring	20, 50, 100, 200, 500
Crossover method	Implementation of selection and reproduction	Elitist Recombination, with probability $p_c=1$
Mutation	Random variation of solutions	Not applied
Stop criterion	Criterion to stop execution of the genetic algorithm	Time step = 500
Migration rate	Probability of selection of solutions for migration	1%, 2%, 5%, 10%, 20%, 50%, 100%
Migration interval	Determines how often migration takes place	Every 1, 2, 5, 10, 20, 50, 100 time steps
Migration type	Selection and copying/moving of solutions for migration	'Move random', 'move best' and 'copy best'

Table 2-2. Parameters and their settings, used in simulations

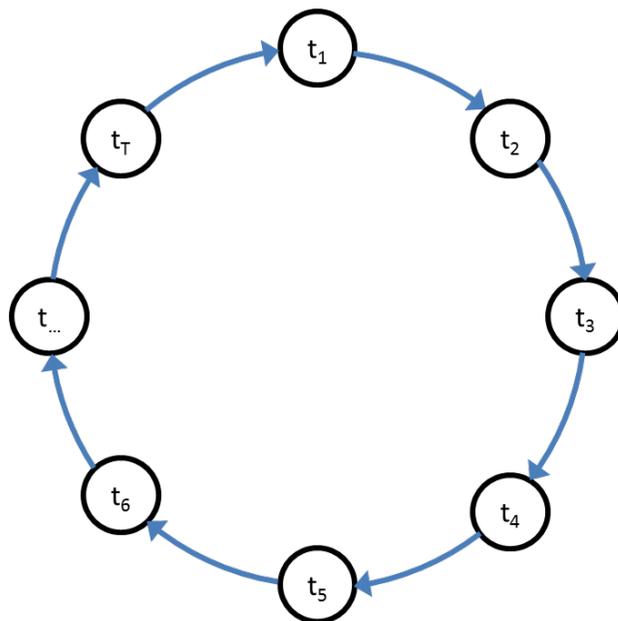


Figure 2-3. Connections between teams in simulations: each team is unidirectionally connected to the next team, the last team is connected to the first team.

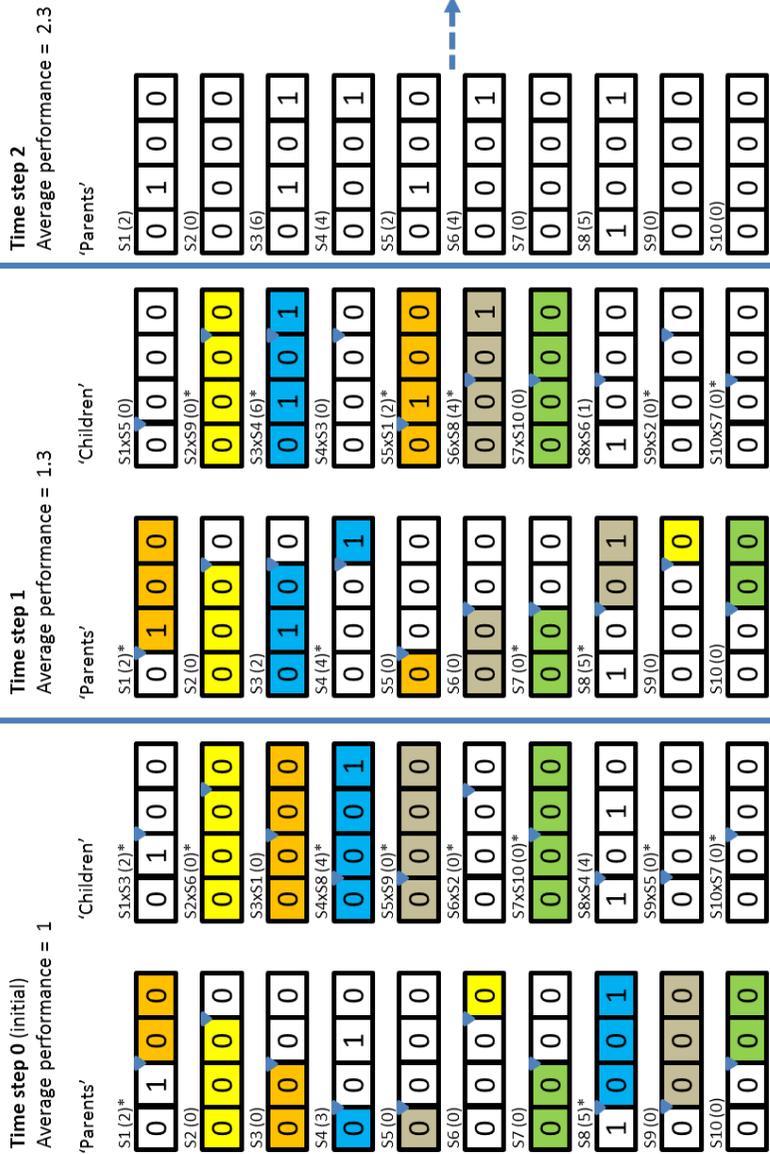


Figure 2-4. Example of the first 3 time steps in a collective approach simulation, with $T=1$, $P=10$ and $N=4$. Performance of individual candidate solutions between (), with blue triangles indicating crossover points and colors indicating parent-child relations. Asterisks* indicate candidate solutions that are selected for the next time step, following the elitist recombination method.

overview of these and other parameter settings that we used in our simulations.

As mentioned before, the network structure of links between firms lies outside the scope of current work, but is addressed in other work in this thesis, mostly in chapters 3 and 4. Here, we interconnect teams, labeled $t_1 \dots t_r$, like a chain of teams, with each team t_i being connected to the next team t_{i+1} with the last team t_r being connected again to the first team t_1 . Each link (or *edge*) is unidirectional (or *directed*): team t_i is connected to team t_{i+1} , but team t_{i+1} is not connected to team t_i . We visually arrange the resulting network of teams in a circular shape, although obviously, this could be any closed loop shape. See figure 2-3.

As a baseline to assess results of the different parallel approaches that were simulated, a collective approach to innovation was simulated, using a serial genetic algorithm. This can be thought of as a single team in which all available resources work together on a single and shared solution set. As indicated in table 2-2, simulations were run for the serial version with a solution set P of 100 bitstrings, all with the same fixed length N per simulation run, having runs for N=20, 50, 100, 200 and 500. Initially, the solution set is filled with bitstrings that were randomly generated, but with a maximum of 10% of random bitstring positions being filled with 1s, the other 90% filled with 0s.

Starting with the initial random population set, a new population set is produced by applying crossover with 'elitist recombination'. For a solution set P = 100, crossover is applied 50 times. For each crossover, 2 parents are randomly selected, producing 2 children – of those 4 candidate solutions, only the two best performing are included in the new solution set, the other 2 are discarded. As explained before, the performance of each bitstring is proportional to the number of 1s, with a weight per 1 depending on its position in the bitstring – the more to the right, the heavier the weight. In terms of an innovating team, crossover on bitstrings or 'genotypes' can be thought of combining elements of different candidate solutions or 'phenotypes', and then selecting those solutions that contain more of the required elements, with a preference for solutions that have more of the more important elements.

This process of generating a new population set is repeated for 500 times, with each new population having an average performance level that is equal to or higher than the former set. As a result, bitstrings will start to contain more 1s, especially on the right – the innovation team is working more and more on candidate solutions that contain more of the required elements, with priority on the most important elements. However, at the same time, crossover will lead to convergence – because of the repeated combination of bitstring fragments, candidate solutions will become ever more similar. This convergence towards a homogeneous solution set with all bitstrings being similar, may happen earlier with a smaller population set size P, but with a speed that also depends on variations that occur in the random process of set initialization, crossover parent selection and migration, with the latter only in the parallel approach.

Having established a performance baseline for the collective approach with a single innovation team, the available resources are then split into a number of teams (10), simulating the parallel approach. So instead of having a single team working on 100 candidate solutions, simulations are run with 10 teams with each team working on 10 candidate solutions. For each team, the process is the same as in the case of the

collective approach, with the important exception that migration is applied. This is done by periodically (following the migration interval parameter) selecting a number of bitstrings (with the number depending on the migration rate parameter and the selection depending in the migration type parameter: selection is either random or based on best performance). The selected bitstrings or 'emigrants' are then sent to the 'next' team where emigrants become 'immigrants', following the circular connectivity topology of the teams. Depending on the migration type parameter, the receiving team may place immigrants in the empty spots of emigrants that left ('move best' or 'move random') or replace randomly selected existing bitstrings by immigrants ('copy best'). Either way, the existing solution set of each team is altered by the effects of migration – new candidate solutions have arrived, existing ones may have left or been replaced. By applying crossover in the next time step, the new knowledge that is represented in new candidate solutions is then recombined with existing knowledge in existing candidate solutions. This allows the innovation team to find new elements that before they didn't know about or had no resources available to address. Importantly, it may also shake up a team that was already converging to a more homogenous but possibly suboptimal solution set – so called 'local convergence'. Gradually, however, knowledge sharing between teams, if repeated long enough, will lead to increasingly similar solution sets across teams, with ultimately one dominant solution being adopted by all teams – so called 'global convergence'.

2.3 Results

To interpret results of our simulations we compare a serial genetic algorithm with several variants of a parallel genetic algorithm. Our comparison is based on the performance that was achieved by each variant at the end of each simulation, averaging 10 different runs for each parameter combination.

First, it can be observed that in almost all of our simulations, the parallel variant outperforms the serial variant, with the exception of cases where a) migration rates are very low (<5%) in combination with migration intervals >2, b) migration intervals are long (>15 for migration rate >5%, > 20 for migration rate > 10%)) and c) migration intervals are short (<5) and migration type is 'copy best', the more so with higher migration rates. Apparently, in cases like a) and b), migration is so weak that teams operate very isolated, to the extent that performance falls below the serial approach. In cases like c), migration is very strong, occurring more often and in greater volume, which combined with migration type 'copy best' makes convergence occur even quicker and more prematurely than in case of the serial approach. See Figure 2-5.

Importantly, in almost all parallel variants, we observe an inverted U-shape for performance along the migration interval dimension (figure 2-4) and along the migration rate dimension (figure 2-6), with the exception of a) sparse migration when the migration rate is very low (<2%), in which case longer migration intervals are always disadvantageous and b) abundant migration when migration intervals are very short, in which case higher migration rates are always disadvantageous. In most cases however, along the migration interval dimension, performance is optimal for mostly intermediate migration intervals (every 5-10 time steps), performance

is poorest for short migration intervals (too little isolation) and poorer for longer migration intervals (too much isolation).

Along the migration rate dimension, it seems that migration rates can increase significantly before performance starts to drop. In the case of migration type 'move random' this drop even seems to be absent. Apparently, migrating randomly selected solutions isn't causing significant negative effects. And also in case of selecting the best individuals for migration, rates can grow > 50% before negative performance effects occur. Moreover, this pattern of a horizontally long inverted U-shape, seems to be quite consistent with increasing migration intervals.

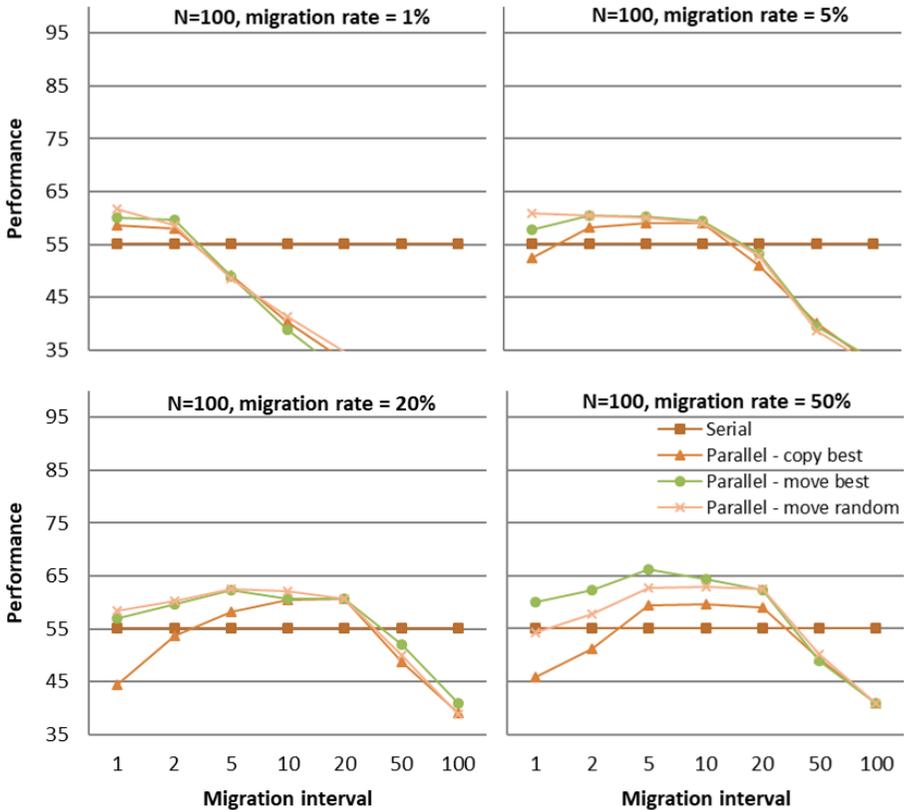


Figure 2-5. Performance comparison for N=100, migration rates 1, 5, 20 and 50%

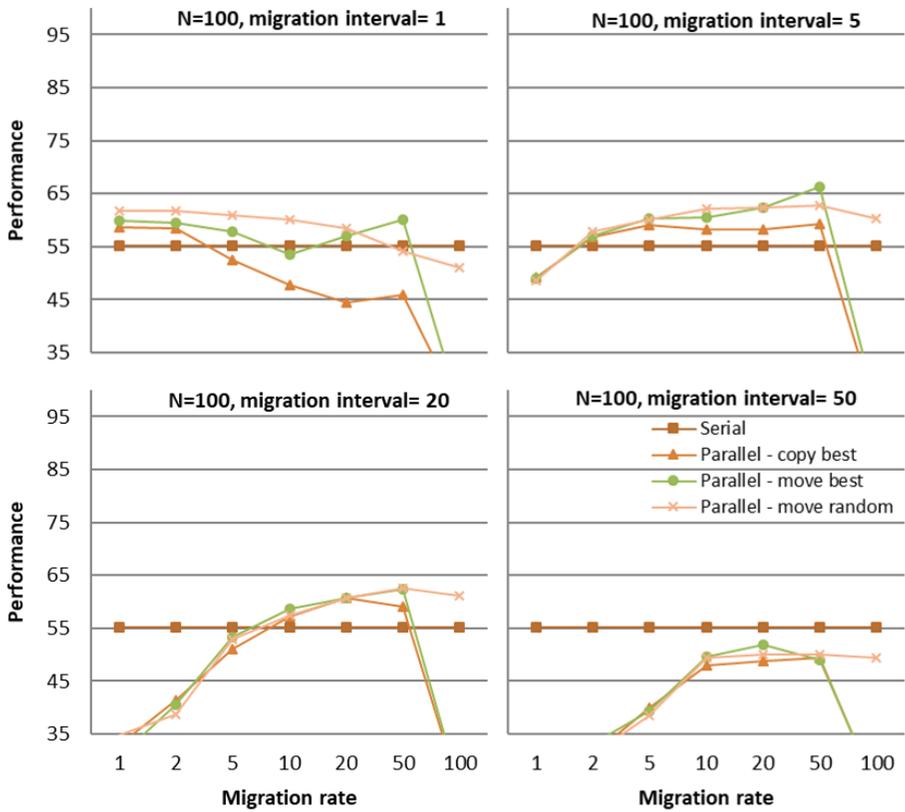


Figure 2-6. Performance comparison for N=100, migration intervals 1, 5, 20 and 50

When examining the interactions between migration rate and interval, we observe that higher migration rates seem to compensate for decreasing performance loss for longer migration intervals. Likewise, shorter migration intervals seem to compensate for smaller migration rates. In other words, migrating less often can be compensated by migrating more, and migrating less can be compensated by migrating more often.

Comparing migration types reveals that differences between migration types seem to increase with higher migration rates and shorter migration intervals. In other words, the impact of migration type is strongest in cases of abundant migration. In such circumstances, copy-best – adding significantly to the speed of convergence – leads to poorest performance, whereas move-best and move-random perform less poor, with more subtle differences between them. We also note that in all other cases, differences between migration types seem to matter less.

With increasing problem difficulty, see figure 2-7, performance goes down. This is a direct result of the fact that in simulations the size of solution sets was kept constant, 100 in total, 10 for each team. In fact, for our most difficult problems (N=200, N=500), performance was more affected negatively by solution set size limitations – as a

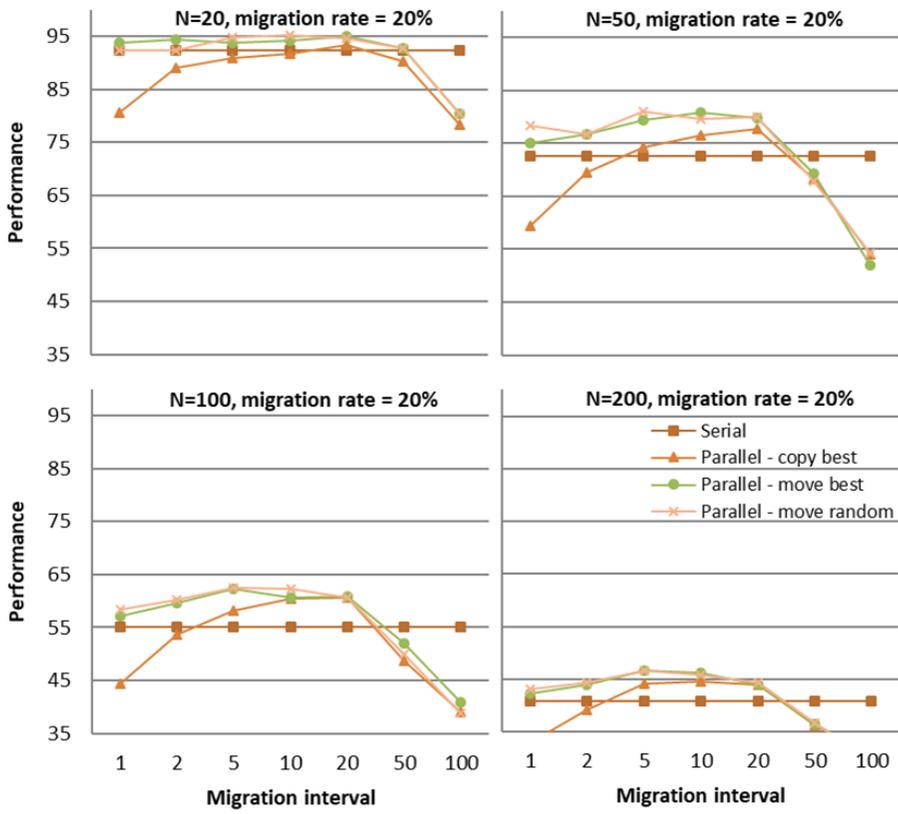


Figure 2-7. Performance comparison for N=20, 50, 100 and 200, for migration rate 20%. Please note that in simulations, solution size remained constant for all N.

result, difference between parallel variants and serial variant became less prominent, although still following the same inverted U shape pattern as at easier problem levels.

Considering different levels of problem difficulty at N=20, 50, 100 and 200, see figure 2-7, we observe that migration type ‘copy-best’ often offers less performance than the other types. As expected, differences between migration types and the serial variant decrease with easier problems, which is most visible in the case of N=20: the problem has become so easy that differences between variants matter less, but again with migration type ‘copy best’ underperforming, even compared to the serial variant, mostly.

On another note, we found that less performing variants in general tended to have bitstrings in which 1s were more prominent to the right side. Of course, this is an inherent result of the ‘counting linear weighted ones problem’, but it also demonstrates what is also known as ‘genetic drift’ (Rogers & Prügel-Bennett, 1999). Bits, or ‘genes’ in biological terms, that do not or hardly matter during some stage of evolution, in time tend to disappear from the genepool in case of limited population

size, unless these useless genes happen to be part of the genome of better performing individuals. Later, when these genes do become more important, they may have disappeared if the population is too small or has converged too quickly, to be dominated prematurely by solutions that have lost these genes in an earlier phase. In other words, in case of the 'counting linear weighted ones problem', the left part of bitstrings typically becomes more important when overall performance is already high (many 1s already to the right side), and when bits on the left side start to make a more significant difference in reaching higher performance levels. However, if these 1s on the left side are no longer present in any solution within the team or collection of teams, as a result of genetic drift caused by premature convergence, then these 1s will never be found and performance will not improve any more.

Our results demonstrate that there are also significant differences between variants when it comes to time. When end results of different variants, as shown above, are at more or less equal performance level, less isolated variants that converge more quickly, offer the benefits of reaching the final performance level earlier. This illustrates the importance of balancing a high speed of convergence on the one hand (achieving innovation results in a shorter time span) and preventing premature convergence on the other hand (achieving better innovation results).

Considering all results, it is clear that all migration parameters are interacting and together set the level of isolation between teams, with best performance results at intermediate levels of isolation, confirming our expectations.

2.4 Conclusion

One of the interesting results of our modeling exercise to deepen our understanding of knowledge sharing between innovating teams, is that the parallel approach seems to be quite robust when it comes to outperforming the collective or serial approach, in line with other work (Cantu-Paz, 1995). Rates can increase to as high as 50% and intervals can increase to as high as 40 time steps, before performance drops below the serial approach. Within those limits, the same holds for all sharing or migration types, with serial only better in case of copy-best and short intervals.

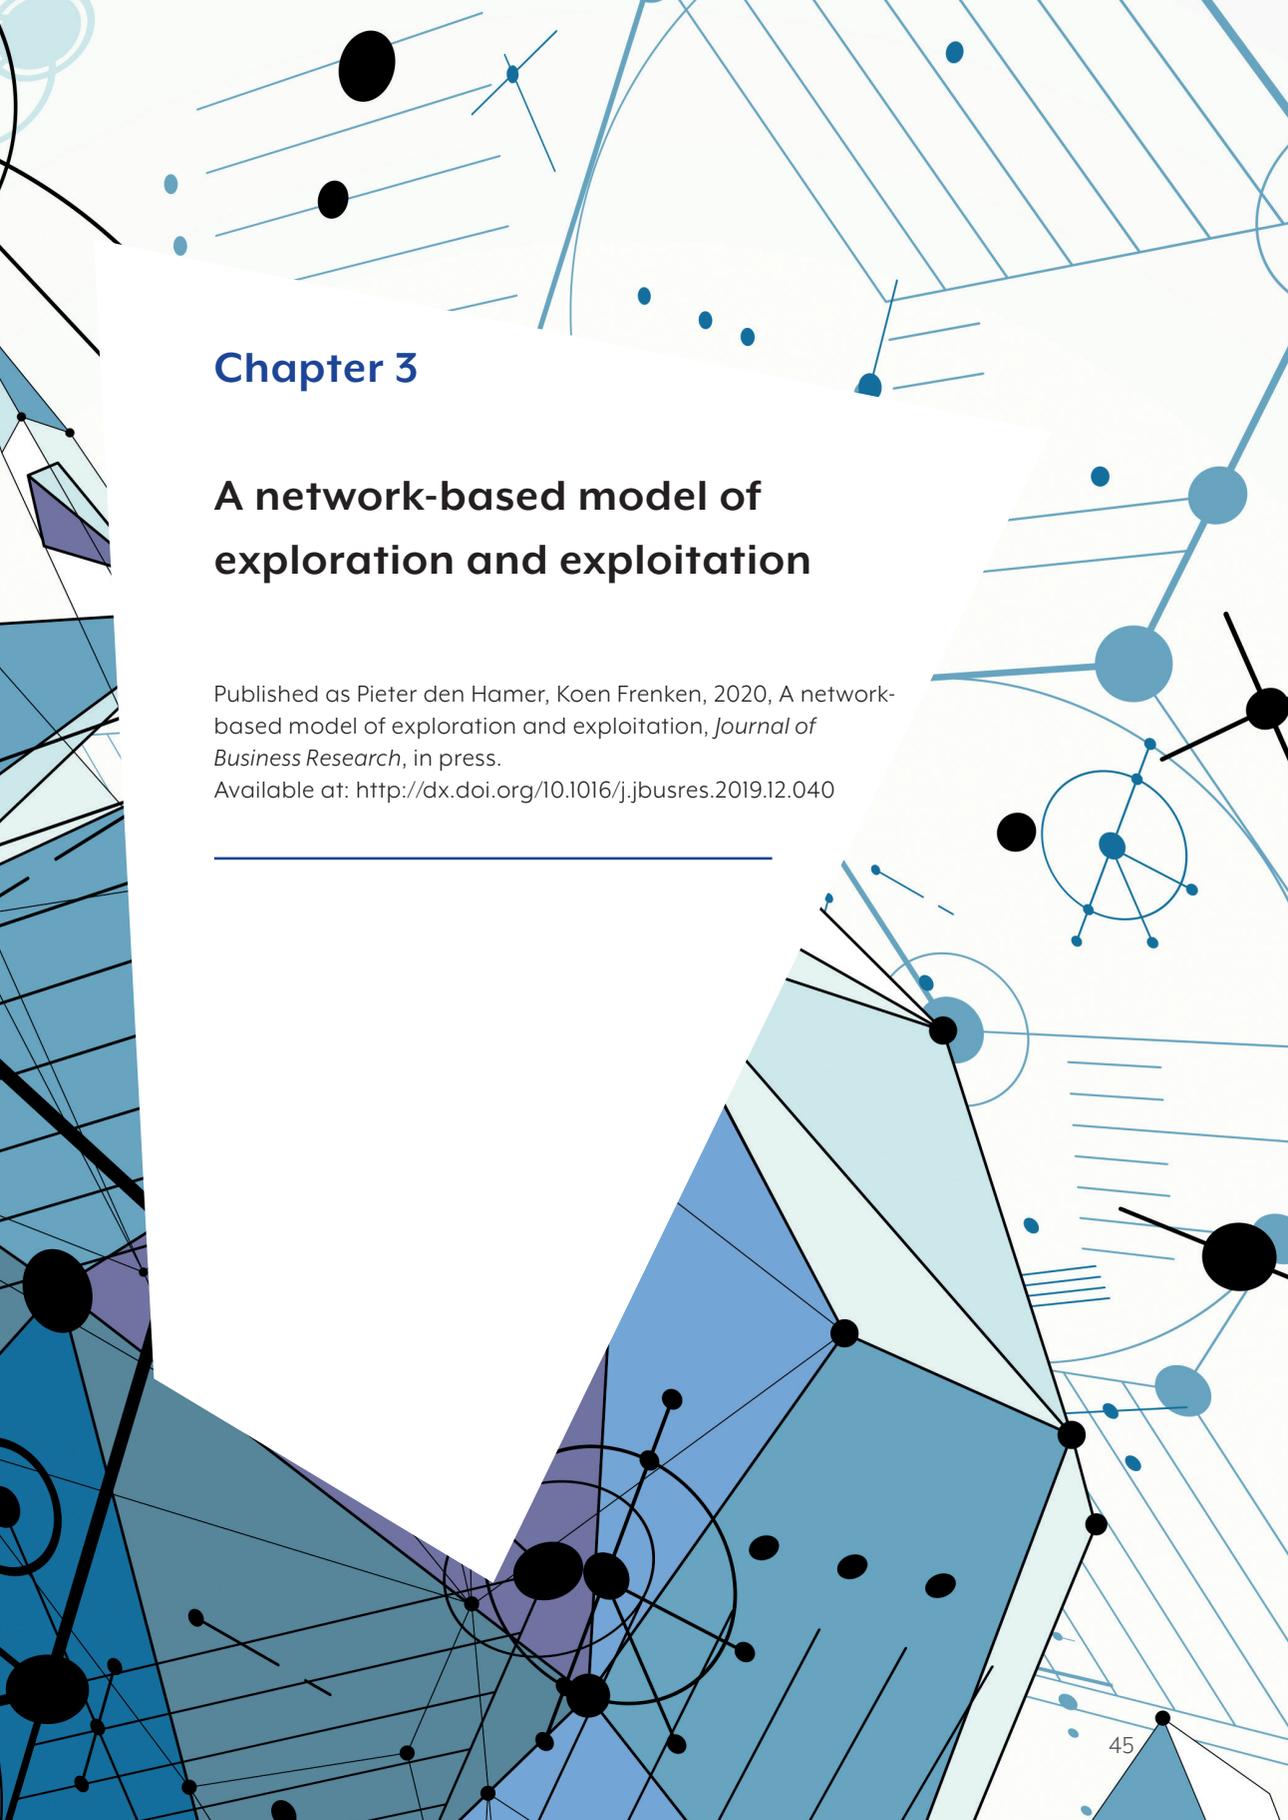
Also, importantly, differences between the serial and parallel approaches and variants thereof become more substantial with a growing difficulty of the problem at hand. When the problem is easier, almost all approaches and variants perform well, and performance differences are small. When the problem is difficult enough, it really starts paying off to fine-tune the parameters that control knowledge sharing, or migration in our model, to reach optimal performance levels. This entails balancing too much sharing on the one hand (Hansen, 1999), leading to premature convergence, with too little sharing on the other hand (Cowan et al., 2014; Gubbins et al., 2014; Nootboom et al., 2007), leading to isolated teams that are too small to solve the problem on their own. To find this balance and optimize knowledge sharing, results in our current work indicate that performance is most sensitive to interval changes, followed by rate and type. However, optimal rate, interval and type are heavily interdependent.

It should be noted that the above observations only hold when population sizes are big enough – if too small for the problem difficulty at hand, then performance will always be poor, regardless of approach or parameter tuning in variants. In other words, more difficult problems require larger solution set sizes – if solution sets become too small, performance degrades to such a low level that differences between parallel and serial approaches become less relevant.

From an innovation policy & management perspective, we conclude that knowledge sharing parameters deserve explicit attention in terms of monitoring and influencing. If the parallel approach is pursued under guidance of a coordinating body like a firm that organizes an innovation contest (Adamczyk et al., 2012; Terwiesch & Xu, 2006), then it is recommendable to support coordination by for example issuing guidelines, providing guidance and active data gathering about actual knowledge sharing, including frequency, amount and strategies to select, retain or move knowledge. By analyzing such data, a coordinating body may assess the impact of knowledge sharing on innovation and results, and may initiate corrective actions or alter their guidelines and guidance, and so on. If a parallel approach exists without an explicit coordinating body, for example in case of competing firms that all do innovation on comparable products or services, government innovation policies may still exist to promote knowledge sharing, even between competing firms through imitation (Nelson & Winter, 1982; Csaszar & Siggelkow, 2010). Regulating or rather influencing knowledge sharing through government innovation policies, implemented by for example requirements in subsidy schemes, will help to achieve better innovation results, both at market / society level and at the level of individual firms, provided that a right balance is found, as underpinned in the current work.

Relevant to innovation policy & management is also the finding that a parallel approach in most cases is preferable to a serial approach, especially when innovation challenges are more daunting in terms of difficulty. If the innovation challenge requires many resources because of its volume or scale, a serial approach may be preferable, or a parallel approach in which the resources per team are sufficiently large. The latter also indicates a possible subject for future work: understanding how, in a parallel approach, performance is impacted by varying the number of teams and their available resources (islands and solution set sizes in our model), while keeping the total amount of resources at the same level. What yields the best performance result, under which conditions: fewer but larger teams that may explore greater parts of their search space on their own, but with less alternative knowledge sources, or many but smaller teams, that can explore less by themselves, but have more alternative knowledge sources (Terwiesch & Xu, 2006)? Answering this question may provide further valuable insights, especially for innovation ecosystems where the number and size of innovating organizations or units can be influenced or even controlled by managers or policy makers.





Chapter 3

A network-based model of exploration and exploitation

Published as Pieter den Hamer, Koen Frenken, 2020, A network-based model of exploration and exploitation, *Journal of Business Research*, in press.

Available at: <http://dx.doi.org/10.1016/j.jbusres.2019.12.040>

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 compared to less complex industries, and that small-world networks yield the highest benefits for collective learning.

3.1 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 (March, 1991; Lavie et al., 2010). 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 et al., 2006; Wilden et al., 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 (March, 1991; Benner & Tushman, 2003). By contrast, exploration involves a search for new and distant knowledge through recombination, experimentation and risk-taking (March, 1991; Savino et al., 2018).

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 et al., 2008).

It is common to assume that imitation is not a blind process, but biased towards successful competitors (Nelson & Winter, 1982; Lieberman & Asaba, 2006). 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 et al., 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 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 et al., 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, albeit stylized, model. We look at firms in different industry contexts characterized with different levels of complexity, as 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 et al., 1995; Powell et al., 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 1) or indirect ties (at social distance > 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 (Miller et al., 2006; Lazer & Friedman, 2007) and related models on inter-firm learning (Cowan et al., 2004; Cowan & Jonard, 2003, 2004, 2009). 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 which knowledge exchange occurs informally.

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). The small-world parameter tunes networks from fully regular to 'small-world' to fully random. The small-world network combines the feature of high clustering of regular network with that of short distances of random networks. Comparing these networks allows us to investigate whether high clustering and short distances are complementary in imitation, as found in alternative models where firms only learn from direct partners (Cowan & Jonard, 2003; Cowan & Jonard, 2004) and in empirical studies on the role of networks on innovation (Capaldo, 2007; Fleming et al., 2007a; Schilling & Phelps, 2007). Second, at the level of dyads concerning each pair of firms, we investigate whether successful instances of learning between two firms occur at an intermediate cognitive distance as implied by the thesis of optimal cognitive proximity (Nooteboom, 2000; Cowan & Jonard, 2009). 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 theoretical 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.

3.2 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 solution 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 (Rikvin, 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 (Rikvin, 2000). The complexity of a product, a technology or service can be thought of in terms of the extent of interdependencies between the components that make up a product, technology or service (Simon, 1962; 1996; Levinthal, 1997; Gatti et al., 2015). 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. That is, 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; Solis-Molina et al., 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). 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) (Lissoni, 2001). 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). What is more, mutual sharing practices are reinforced by professional and academic norms in communities of practices (Lissoni, 2001).

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 (Powell et al., 1996; Pyka, 2002; Breschi & Lissoni, 2009). 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: Ozman, 2009; Phelps et al., 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'), but also on its indirect ties ('friends of friends') (Ahuja, 2000; Breschi & Lissoni, 2009). In general, one can expect that knowledge from socially proximate firms is more accessible than knowledge from more 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 (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 et al., 2007a; 2007b; 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 needed 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* (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 and 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.

3.3 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 much less complex (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 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 them 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.

3.3.1 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 for 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 Ganco and Hoetker (2009), Puranam et al. (2015) and Baumann et al. (2019).

In the NK-model, an organization'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

Assuming that firms exploit their knowledge base through local 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).

3.3.2 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 (for example, 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 (for example, technical efficiency of a technology) and not to the economic value of the firm as such (for example, 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 make 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 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. This is 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 of 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 making 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 bit that a firm attempts 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 between firms D (called ‘degree’) 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 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 ‘short-cuts’ (Figure 3-2). When further rewiring all ties as to obtain a random network, the clustering of nodes 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 et al., 2007a; Lissoni et al., 2013; Newman, 2003; Uzzi et al., 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.

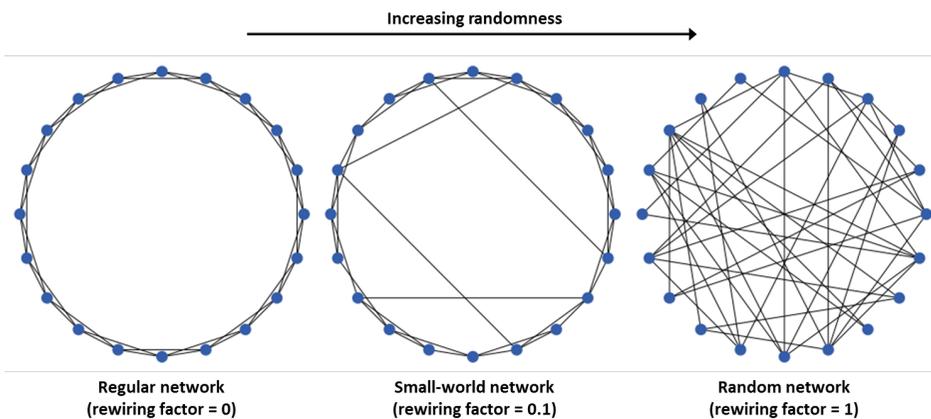


Figure 3-2. Three examples of networks of firms ($F=20, D=4$). Regular networks have high clustering. Networks are created by following the method of Watts and Strogatz (1998), varying the ‘rewiring factor’ between 0 and 1.

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]

Table 3-1. Parameter settings and output variables

3.3.3 Simulation

In each simulation, firms start from randomly assigned bitstrings and engage in search for alternating periods of exploitation and exploration. Exploitation consist of local search until a local peak is reached. Once a local peak is reached, a firm attempts 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 than initially envisaged (Posen et al., 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.² 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 3-1. 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

² 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.

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:

```

;initiation
random-select-network ( $\beta$  parameter, F=100)
 $\Delta$  = 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,  $\Delta$ , cd)
        if performance(b-new) > performance(b-current) then
            bitstring(f) = b-new
        end if
    next firm f
next time-step

```

3.4 Results

3.4.1 Network-level

Figure 3-3 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

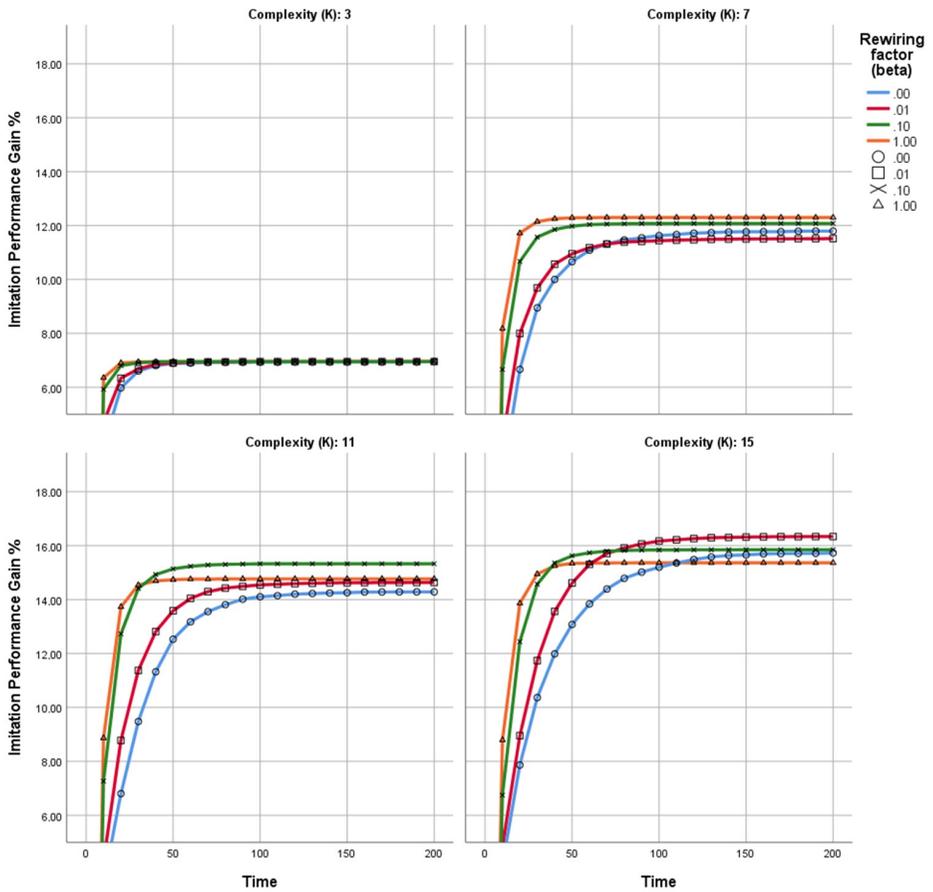


Figure 3-3. Performance gain achieved by exploration compared to the baseline with firms only engaging in exploitation. Averages over 100 simulation runs.

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 Figure 3-3 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 firms 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 having large basins of attraction. 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 will then 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 $K=15$.³ 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 Figure 3-3 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 more subtle effect. One 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 Figure 3-3, 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 ($K=15$).

3.4.2 Dyad-level

Figure 3-4 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 attempts 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.

3 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 $K=3$.

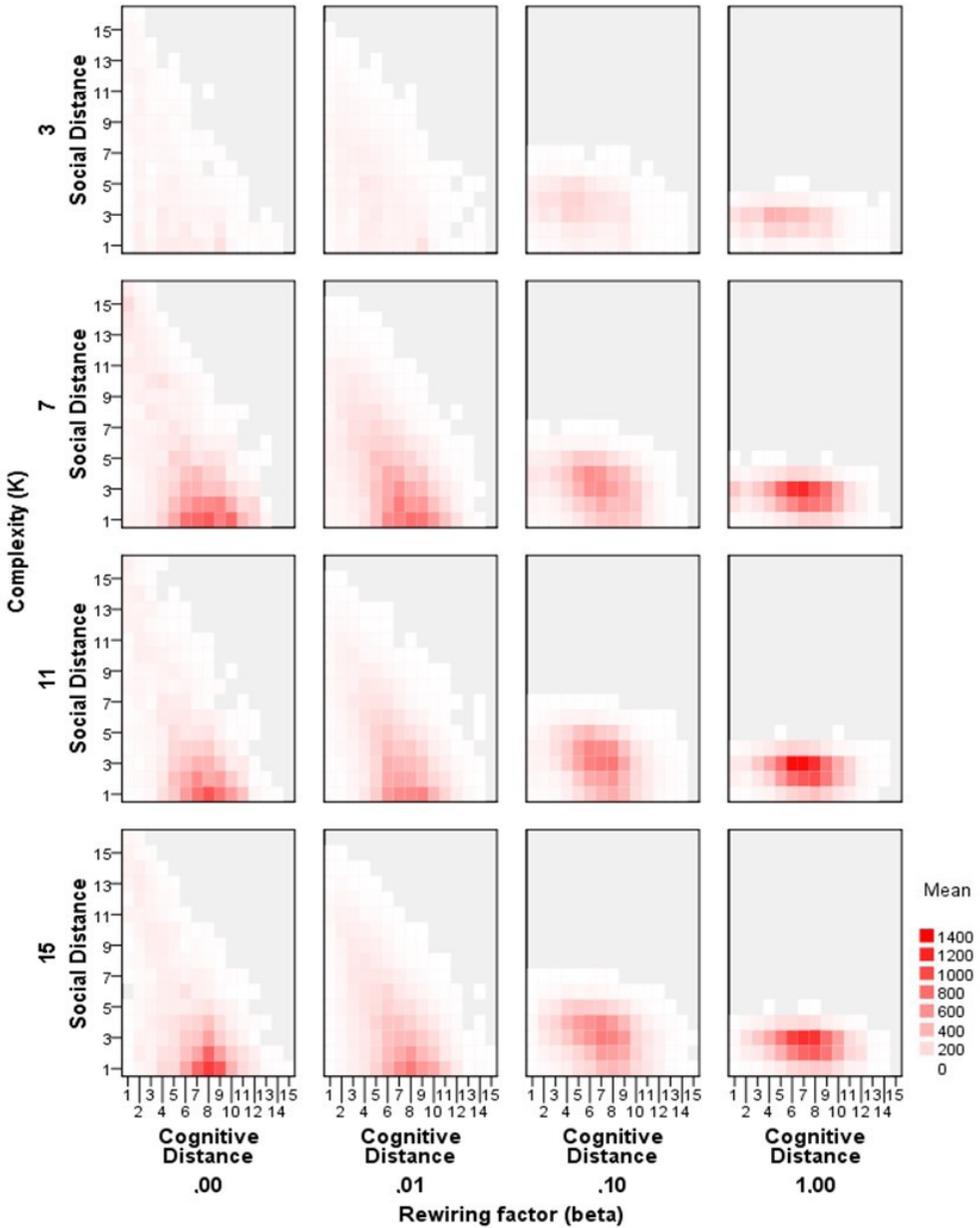


Figure 3-4. Occurrence (darker color corresponds with higher occurrence) of successful imitation, for different complexity (K) levels (rows) and rewiring factors (β) (columns), at different cognitive and social distances (reverse of proximity). Means taken over 100 simulation runs. For lower complexity ($K=3$) 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.

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 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 (Wuyts et al., 2005; Nooteboom et al., 2007; Gilsing et al., 2008; Fitjar et al., 2016).

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 located at a distance in the social network, 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* (Huber, 2012; Cassi & Plunket, 2015; Steinmo & Rasmussen, 2016).

3.4.3 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 finding a high peak will increase (decrease).

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 of firms 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.

3.5 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 and 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 previous 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: (i) the higher value of exploration in highly complex industries compared to less complex industries, and (ii), the benefits of small-world networks in collective learning.

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 and from a large range 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 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 (Levthinal, 1997). Then, a question to address is what optimal balance exists between internal and external exploration depending on a firm's

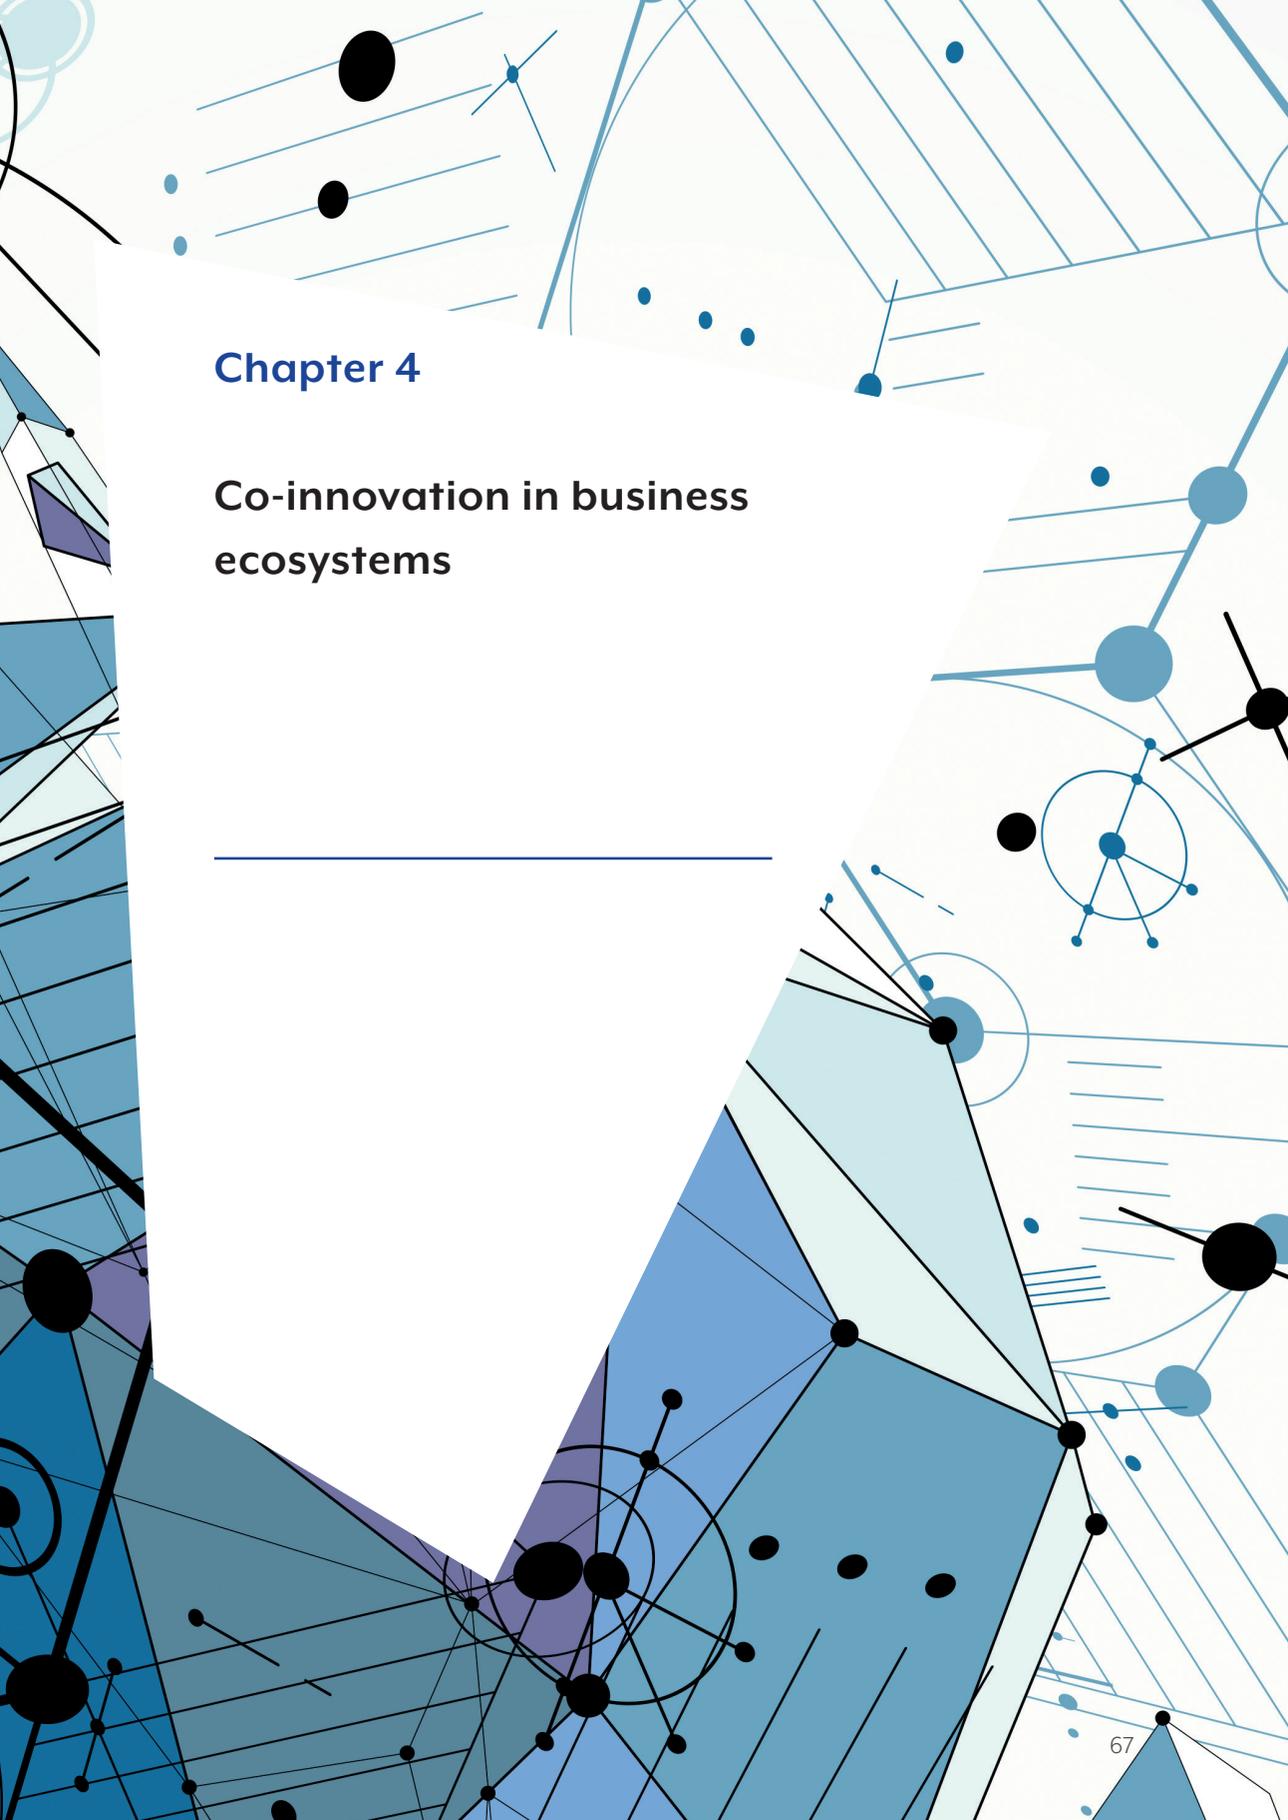
competitive position. In particular, well-performing firms can arguably learn more from exploration by internal search compared to poorly performing firms that can benefit more from imitating others. A second extension is to incorporate cost. In our model, we abstracted away from the cost of imitation. 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 context of high product complexity. Hence, strategically maneuvering into a favorable network position combining short distances with high clustering only matters in context where products, and the underlying knowledge base, is complex. Second, firms should refrain from attempting to learn from 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 specially 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 structures acquires small-world characteristics. Indeed, as a public actor, government can influence the macro-level structures of collaboration among firms by subsidizing inter-firm networks (Van Rijnsoever et al., 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 increasing 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.



The background of the page is a complex, abstract geometric design. It features a variety of blue and black lines, circles, and polygons. Some lines are straight and parallel, while others are curved or intersecting. There are several large black circles and smaller blue circles scattered throughout. The overall effect is that of a technical or scientific diagram, possibly representing a network or a system of interconnected parts. The colors are primarily light blue, dark blue, and black, set against a white background.

Chapter 4

Co-innovation in business ecosystems

Abstract

Innovations by one firm often trigger innovations in other firms, leading to cascades of innovations. We model such co-innovation patterns in business ecosystems by applying Kauffman's NKCS-model from theoretical biology. We extend the NKCS-model in that we vary the ecosystem network structure from regular to small-world to random. We further investigate the resilience of ecosystem firms against shocks. Findings indicate that the ecosystem's network structure is a relevant aspect in the ecosystem's regime. Network structure, together with the level of internal and external complexity, determine if an ecosystem is either in a rigid, chaotic, or in-between regime, known as 'edge of chaos' (EoC). In EoC regime, balancing stability with adaptability, innovation performance levels are highest, with clustered innovations occurring and innovation cascade sizes following a power law distribution, confirming and underpinning known stylized facts about (co)innovation and empirical work.

4.1 Introduction

Innovation does not happen in isolation. Innovations by one firm are often triggered by innovations by other firms, be it from suppliers, clients, competitors or partners. Together, firms operate in a business ecosystem where firms interact as suppliers, clients, partners or competitors, or even in multiple context-dependent roles at the same time. In ecosystems, firms 'co-innovate' (Adner, 2012): innovation in one firm goes hand in hand with innovation in one or more other firms, either coordinated and willingly, like clients and suppliers working together in their innovation, or uncoordinated or even unwillingly, like innovation that is required to keep up with new customer requirements or competitor moves.

There are many examples of how firms, the market and society at large benefit from firms doing co-innovation. Apple's iPhone, as one of the first smart phones on the market, is really the result of many other innovations made by other firms and organizations, including touch screen, battery, optical and other technologies (Lazonick et al., 2013). New innovations in these areas may trigger the development of new smart phone models, and *vice versa*. In other cases, co-innovation takes place with strong involvement of customers, which is also known as 'co-creation' (Sawhney et al., 2005). A well-known example is the case of Lego, a firm that is actively working with their young and old customers to develop new models of their toy building kits, which in turn may lead to new requirements for Lego's suppliers, triggering innovations like for example sensor and actuator building blocks for Lego robotics products.

Innovations that trigger other innovations may trigger yet other innovations, and so on, causing a cascade or time-clustered series of innovation (Schumpeter, 1939; Silverberg & Verspagen, 2005; Arthur, 2007). That is, innovations do not occur randomly in time, but occur punctuated and clustered in certain periods. In this context, Schumpeter made a distinction between radical and incremental innovations (Schumpeter, 1939), with the latter following up the former. Radical innovations open the gate or provide a break-through to a whole new area and era of innovation opportunities. Dosi (1988), based on his analysis of patent data, confirmed that innovation clusters or cascades vary in size, with most being very small and few very large. Consequently, he reasoned that innovation cluster sizes may follow a power law distribution where small events are much more common than large events, but still according to a very regular pattern of scaling. A classic example is the cluster of innovations in the 19th century to harness electricity, in turn triggering a burst of innovations in which electricity was applied in transportation, domestic appliances and many other domains (Arthur, 2007). Clusters of innovations following a radical innovation have also been associated with a paradigm shift or 'socio-technical transition', in the sense of a large-scale change that occurs suddenly yet endogenously (Frenken et al., 2012).

While the Schumpeterian notions of radical and incremental innovation, and related concepts like paradigms and transitions, are central to innovation thinking for the past four decades, only few attempts have been made to model such innovation patterns from an ecosystem perspective. A more fundamental understanding may come from an evolutionary and complexity perspective, in which we take the

business ecosystem quite literally as an ecosystem in which firms interact according to a certain network structure (Kauffman & MacReady, 1995). Firms are born, live and die, continuously adapting to their environment of customers, suppliers, competitors and general market conditions. Firms adapt by evolving their structure, processes, products and services. In this perspective, innovation is a response to opportunities and threats, like new innovations made by other firms, which may be used to improve a firm's own products or services or to increase efficiency of existing processes. But, as we will argue, the impact on one firm of innovations made by other firms, is very much dependent on both the structure and the extent of dependencies between firms, and on the extent to which firms are actually able to respond to changes in their environment (Kauffman & MacReady, 1995; Levitan et al., 2002; Press 2008).

In this study, we develop an ecosystem model of co-innovation to gain a further understanding of the ecosystem nature of innovation and stylized facts about clustering of innovation (Arthur, 2007; Verspagen & Duysters, 2004; Dosi, 1988). We do so by applying and extending the NKCS-model from theoretical biology (Kauffman & Johnsen, 1991), which has been applied before in economics and management science (Vidgen & Bull, 2011). Our main contribution is to investigate the role of the network structure that characterizes the business ecosystem and, more specifically, to test whether a 'small world' structure – often empirically observed (Fleming et al., 2007a; Lissoni et al., 2013; Newman, 2003; Uzzi et al., 2007) – provides the best collective performance and resilience. We further show how a business ecosystem can evolve endogenously towards such a small-world network structure. We conclude by identifying implications for innovation policy and management.

4.2 Background

Ecosystems (Tansley, 1935) in general are considered to be systems that comprise the whole of interacting actors and their environment. Like in biological ecosystems, the network of interactions between firms in a business ecosystem is an important factor in understanding how actors like organisms or firms respond to changes in their environment or in other actors (Kauffman & Johnsen, 1991; Vidgen & Bull, 2011). Organisms and firms do not evolve in isolation. They co-evolve, continuously responding to changes in their ecosystem which may have an impact on their own fitness or performance. For example, if a predator species happens to evolve in such a way that increases its speed in hunting, this will create a selective pressure on prey species, favoring those that are quickest to escape, which in turn may have an impact of food sources of the prey, which may be other animals or plants. A key observation here, in the case of this highly simplified example, is that the network structure of the ecosystem is an important factor in the extent to which a change in one species has an impact on other species. If the predator in this example hunts only for one type of prey, which in turn only eats one type of food, then the impact of the initial change might be limited. In that case, the network is very sparse and limits the propagation of change. In the other extreme, where a predator has a broad choice of prey species, which in turn have a big and diverse menu to fulfill their food needs, and so on, the network is very dense, supporting the diffusion of



a1	a2	a3	b1	b2	b3	p(a1)	p(a2)	p(a3)	p(A)	b1	b2	b3	a1	a2	a3	p(b1)	p(b2)	p(b3)	p(B)
0	0	0	0	0	0	0.12	0.04	0.95	0.37	0	0	0	0	0	0	0.23	0.65	0.08	0.32
0	0	0	0	1	0	0.12	0.20	0.95	0.42	0	0	0	0	1	0	0.78	0.65	0.75	0.73
0	0	0	1	0	0	0.34	0.04	0.06	0.15	0	0	0	1	0	0	0.23	0.43	0.08	0.25
0	0	0	1	1	0	0.34	0.20	0.06	0.20	0	0	0	1	1	0	0.78	0.43	0.75	0.65
0	0	1	0	0	0	0.12	0.86	0.26	0.41	0	0	1	0	0	0	0.23	0.65	0.60	0.49
0	0	1	0	1	0	0.12	0.08	0.26	0.15	0	0	1	0	1	0	0.78	0.65	0.88	0.77
0	0	1	1	0	0	0.34	0.86	0.59	0.60	0	0	1	1	0	0	0.23	0.43	0.60	0.42
0	0	1	1	1	0	0.34	0.08	0.59	0.34	0	0	1	1	1	0	0.78	0.43	0.88	0.70
0	1	0	0	0	0	0.13	0.62	0.45	0.40	0	1	0	0	0	0	0.65	0.32	0.08	0.35
0	1	0	0	1	0	0.13	0.10	0.45	0.23	0	1	0	0	1	0	0.14	0.32	0.75	0.40
0	1	0	1	0	0	0.67	0.62	0.30	0.53	0	1	0	1	0	0	0.65	0.77	0.08	0.50
0	1	0	1	1	0	0.67	0.10	0.30	0.36	0	1	0	1	1	0	0.14	0.77	0.75	0.55
0	1	1	0	0	0	0.13	0.89	0.52	0.51	0	1	1	0	0	0	0.65	0.32	0.60	0.52
0	1	1	0	1	0	0.13	0.77	0.52	0.47	0	1	1	0	1	0	0.14	0.32	0.88	0.45
0	1	1	1	0	0	0.67	0.89	0.98	0.85	0	1	1	1	0	0	0.65	0.77	0.60	0.67
0	1	1	1	1	0	0.67	0.77	0.98	0.81	0	1	1	1	1	0	0.14	0.77	0.88	0.60
1	0	0	0	0	0	0.19	0.04	0.95	0.39	1	0	0	0	0	0	0.55	0.09	0.17	0.27
1	0	0	0	1	0	0.19	0.20	0.95	0.45	1	0	0	0	1	0	0.61	0.09	0.27	0.32
1	0	0	1	0	0	0.38	0.04	0.06	0.16	1	0	0	1	0	0	0.55	0.86	0.17	0.53
1	0	0	1	1	0	0.38	0.20	0.06	0.21	1	0	0	1	1	0	0.61	0.86	0.27	0.58
1	0	1	0	0	0	0.19	0.86	0.26	0.44	1	0	1	0	0	0	0.55	0.09	0.71	0.45
1	0	1	0	1	0	0.19	0.08	0.26	0.18	1	0	1	0	1	0	0.61	0.09	0.42	0.37
1	0	1	1	0	0	0.38	0.86	0.59	0.61	1	0	1	1	0	0	0.55	0.86	0.71	0.71
1	0	1	1	1	0	0.38	0.08	0.59	0.35	1	0	1	1	1	0	0.61	0.86	0.42	0.63
1	1	0	0	0	0	0.63	0.62	0.45	0.57	1	1	0	0	0	0	0.96	0.90	0.17	0.68
1	1	0	0	1	0	0.63	0.10	0.45	0.39	1	1	0	0	1	0	0.22	0.90	0.27	0.46
1	1	0	1	0	0	0.44	0.62	0.30	0.45	1	1	0	1	0	0	0.96	0.49	0.17	0.54
1	1	0	1	1	0	0.44	0.10	0.30	0.28	1	1	0	1	1	0	0.22	0.49	0.27	0.33
1	1	1	0	0	0	0.63	0.89	0.52	0.68	1	1	1	0	0	0	0.96	0.90	0.71	0.86
1	1	1	0	1	0	0.63	0.77	0.52	0.64	1	1	1	0	1	0	0.22	0.90	0.42	0.51
1	1	1	1	0	0	0.44	0.89	0.98	0.77	1	1	1	1	0	0	0.96	0.49	0.71	0.72
1	1	1	1	1	0	0.44	0.77	0.98	0.73	1	1	1	1	1	0	0.22	0.49	0.42	0.38

Figure 4-1. Simple example of a NKCS-model and associated performance table, with $N=3$, $K=1$, $C=1$ and $S=2$. For example component a_1 depends on a_2 and b_1 , resulting in 8 possible configurations for a_1 , for each of which a random performance value is drawn. Note that in this particular example, firm A depends on 2 components of firm B, resulting in 4 different performance values for each total configuration of firm A. So even if firm A remains unchanged as for example '0 0 0', the performance of firm A will be affected as a result of changes in firm B – in effect deforming A's performance landscape.

the change (Newman, 1999). Similarly, it is expected that ecosystem networks with a more clustered structure, will result in changes or cascades of changes that remain more within a cluster, whereas more random networks will facilitate a higher and quicker diffusion, similar to epidemic disease spreading (Newman, 1999). And also in the context of innovation, there is growing consensus and empirical evidence that the characteristics of firm networks and the position of firms within networks, are relevant to innovation performance (for a review, see Ozman, 2009).

Bringing the above example to the context of co-innovation, indicates how one firm's innovation may trigger innovation in other firms. Interestingly, this process may continue and create a cascade of innovations. This may even build up to a 'transition', depending on the size and impact of the cascade. It is even possible that the firm that caused the initial trigger innovation, may become itself part of the train of responding innovations or cascade. In that case the firm must respond to innovations in another firm that were caused by its own initial innovation: dependencies between firms may be mutual, either directly or indirectly. But the original firm's response may cause the other firms yet again to innovate, which again may affect the original firm. This process ends when involved firms can no longer find ways to improve their position, at which point a Nash equilibrium (Nash, 1951) has been reached (Kauffman & Johnsen, 1991). But the process may also continue for a long time or indefinitely, causing what is called 'a red queen dance' by biologist Van Valen (1977) inspired by Lewis Carroll's story about Alice in which she finds the Red Queen – the two run as fast as they can, but remain in the same place. This notion gave rise to the Red Queen's hypothesis in evolutionary theory, stating that co-evolving systems need continual development in order to maintain existence (Van Valen, 1977).

Another important factor to consider when it comes to understanding how one innovation may lead to another, is the extent of dependencies between actors like firms or organisms in an ecosystem (Kauffman & MacReady, 1995; Kauffman & Johnsen, 1991; Levitan et al., 2002; Press, 2008). If there is low dependency between actors, changes in one actor will have a more limited impact on the other actor, and *vice versa*. For instance, if a firm depends on a component from one of its suppliers, this component may be highly critical to the quality of the product at hand. In contrast, the component may also be commodity and therefore easily interchangeable with other components from the same or other suppliers. Consequently, innovations in commodity or interchangeable components will create a smaller impact than innovations in highly critical components.

Such dependencies or complexity, may exist *externally*, between components of firm A and components of firm B, but may also exist *internally* between the components of a product of firm A (Simon, 1996). For example in a car, internal complexity exists because changes in one component, an increase in the engine power, also require changes in another component, the brakes, in order to maintain the car's safety and other characteristics. External complexity, in this example, exists because for instance the quality of the brakes depends on the characteristics of the brake cylinder that is provided by one of the car manufacturer's suppliers, which in turn may depend on the characteristics of the steel or composite material that is provided by yet another supplier to the brake cylinder supplier.

Interestingly, it should be noted that internal complexity is actually inhibiting the extent to which a firm may be able to respond to a change in its environment. If internal complexity is high, a change in one component, in response to an external dependency for that component, will require changes in many other components of the same product, as well. And *vice versa*, if internal complexity is low, then changes in one component are more easily accommodated, as the impact on other components of the same product is limited (Kauffman & MacReady, 1995; Levitan et al., 2002).

4.3 The NKCS-model

Co-innovation in the context of a business ecosystem with multiple complementary or competing firms is modeled here by applying the NKCS-model of Kauffman and Johnsen (1991). In this seminal work, the NKCS-model was originally intended to represent coevolution between species in a biological ecosystem, but has also been used in organizational and management science studies to represent business ecosystems from an evolutionary and complex systems perspective, both conceptual (Baum, 1999; Kauffman & MacReady, 1995) and empirical (Marion, 1999; Colovic & Cartier, 2007), with only a few studies based on simulations (Chang & Harrington, 2000; Chang & Harrington, 2003; Levitan et al., 2002). Following his own work and the proposals of Ahouse et al. (1991) on the use of NKCS outside biology, Kauffman himself and MacReady were amongst the first to apply the NKCS-model in the context of technological co-innovation (Kauffman & Macready, 1995). Likewise, Stewart (2001) applied the NKCS model for structuring firm organizations and Baum (1999) identified strategies for firms to manage or influence their internal and external complexity (K and C) levels. Above organizational level, McKelvey (1999), Chang and Harrington (2000; 2003) and Marion (1999) applied NKCS to respectively value chain, retail and microcomputer industry perspectives. Caminati (1999) applied the model to explain technology coevolution between sectors. Celo et al. (2018) analyzed the impact of outsourcing on the performance of global factories by using the NKCS notions of internal and external complexity, revealing that a balance between these must be sought by preventing excessive outsourcing, or *vice versa* (Celo et al., 2018). Other balancing acts, like searching for innovation priorities by taking either internal or external/competitive considerations into account, were also modeled by NKCS (Moran et al., 2011). For a further overview, see (Vidgen & Bull, 2011).

The NKCS-model is an extension of the NK-model (Kauffman, 1993), which in an innovation context is used to represent subjects of innovation (a product, service or process) with N components and K interdependencies between these components, reflecting internal complexity (Frenken, 2006b). If $K = 0$, there are no dependencies between components and the performance of the subject of innovation as a whole is linearly related to the performance of each component. With increasing K (maximum $K = N - 1$), the performance of individual components becomes more dependent on the performance of other components. In other words, a change in one component impacts the performance of one or more other components, thus increasing complexity and making it harder to do find ways of improving the overall performance, as an improvement in one component may lead to degradation in another. Using the landscape metaphor with heights

representing performance levels (Wright, 1932) of all possible configurations: with lowest K , the landscape has a single peak, and with increasing K , the landscape becomes hillier, and with highest K , the landscape has become very rugged. The difficulty of traveling through these different types of landscapes corresponds to the level of complexity and the difficulty of finding optimal innovation results. See chapter 3 for a further explanation of the NK-model.

The NKCS-model extends the NK-model with the notion of having multiple species S – or actors like firms in the current context – in a shared ecosystem with C dependencies between their components. C is similar to K , with the important difference that K represents the number of dependencies between components within an actor, or *internal* complexity, and C represents the number of dependencies between components of two different actors, or *external* complexity. The implication of having external complexity is that changes in one or more components of firm B will impact the performance of components in firm A, or *vice versa*. In other words, changes in firm B will deform the performance landscape of firm A: if A was on a local peak with fairly good performance, after changes in B, A might find itself in a valley, with poorer performance, although its own configuration hasn't changed. See figure 4-1 for a simple example, showing both internal dependencies within firm A and firm B, and external dependencies between firm A and firm B. Performance of firm A changes not only when components within A change, but also when components in B (to which A is dependent) change, even when A itself does not change. In terms of the earlier example: if firm A is a car manufacturer, and firm B is a supplier of say brake cylinders, a drop in the quality of these cylinders will result in poorer car performance, although the car configuration itself hasn't changed.

In the general NKCS-model, external dependencies are mutual between firms, but not necessarily between components. In other words, if firm A depends on firm B, then firm B also depends on firm A. But if component a_1 of firm A depends on b_1 of firm B, then component b_1 does not necessarily depend on a_1 . However, in case of $C=1$, b_1 depends on exactly one component from firm A, either a_1 , a_2 or a_3 , in case of $N=3$. Without loss of generality, the general NKCS-model is employed with each component having 2 alleles, say 0 and 1, which can be thought of as the component being present or absent. With each component depending on $K+C$ other components, it follows that there are 2^{K+C} possible configurations for each component and the components on which it depends. For each configuration, a random performance value p is drawn that lies between 0 and 1, as the generalized and theoretical NKCS-model is agnostic to a specific environment. The performance p for the firm as a whole (or product, service or process thereof, depending on the level of model application) is then simply calculated by taking the average of the performance value of all current component configurations. Figure 4-1 includes a performance table, in this case for $N=3$, $K=1$, $C=1$ and $S=2$, corresponding to the dependencies that are depicted above the table, illustrating the way performances are calculated.

In general, the NKCS-model provides a way of representing a business ecosystem in which firms co-innovate⁴. S may be interpreted as the number of competing/

4 It should be noted that co-innovation has also been modeled with only using the NK-model and other game theoretic approaches on co-innovation (Zeppini et al., 2014). However, in the case of using NK to model co-innovation, the complexity of the subject of innovation in each firm is not taken into account, as K is only used as a general measure to indicate the extent of coupling or dependency between firms. Or, as in the case of the work by Press (2008), the NK model is used with a predefined structure of dependencies between parts and different firms being associated with specific subsets of parts.

complementary firms, C as a measure of the intensity of inter-firm complementarities, K as average complexity per technology with N components.

4.4 Ecosystem regimes in NKCS

One of the important results in the original NKCS work (Kauffman & Johnsen, 1991) holds that ecosystems may have varying degrees of flexibility or adaptiveness. In one extreme, ecosystems are rigid: actors hardly evolve, because making changes to their configuration is difficult without losing performance. Firms in such a regime innovate until they have reached a local peak, and then find themselves stuck at that peak, at which point innovation and learning halts. In another extreme, ecosystems are chaotic: actors innovate all the time, continuously responding to changes in other actors, but fail to improve their fitness. In this regime, firms do not reach any peak, because peaks appear and disappear as their landscape is constantly changing ‘under their feet’.

Importantly, in-between these two states of rigidity and chaos, a transition state was identified, called the ‘edge of chaos’ (Kauffman & Johnsen, 1991). In this state, the ecosystem balances stability with adaptiveness: actors are still able to evolve or learn, thus maintaining or improving performance levels, which are not disrupted by continuous changes in their environment. In landscape terms: peaks remain long enough in one place for firms to climb them, while new features in the landscape, like new nearby peaks, ridges or less steep slopes, allow firms to escape from their current peak to another peak. Consequently, performance levels in the ‘learning’ or ‘edge of chaos’ regime are highest compared to performance that is hardly improving under a rigid regime, or performance that is continuously undermined in a chaotic setting (Kauffman & Johnsen, 1991).

Internal complexity (determined by K, relative to N)	External complexity (determined by C relative to N, S, D and network structure)	Regime	Performance
Low	High	Chaotic	Low
High	Low	Rigid	Low
Balanced	Balanced	Edge of Chaos	High

Table 4-1. Combinations of internal and external complexity and corresponding regimes and performance levels.

As was pointed out by Kauffman & Johnsen (1991), the regime in which an ecosystem finds itself is to a large degree determined by K and C (see table 4-1). Low K makes it relatively easier for firms to adapt to innovations in other firms, or to other external changes – creating more adaptability (Press, 2008), whereas high K will keep firms from responding quickly, if at all, to innovations in other firms, leading to more stability. Low C implies that firms are less dependent on innovations in other firms, creating more stability, whereas high C requires firms to respond to more innovations in other firms, implying

more flexibility. The original NKCS-model by Kauffman & Johnsen (1991) also identified other factors that influence the ecosystem regime, albeit to a lesser extent. Rising the number of firms (S) or increasing the number of firms (D) with whom each firm has complementarities, will both push the ecosystem in a chaotic direction, and *vice versa*. Levitan et al. (2002) refined the original NKCS work to allow finetuning of particularly the number of firms (S) and its effect on stability, flexibility and overall performance.

The key point here is that the performance that firms can reach by innovation, is relatively highest in the edge of chaos regime, where internal and external complexity are balanced. In this regime firms experience sufficient stability to climb a peak in their current landscape and exploit the performance benefits thereof, while still being able to adapt themselves to new circumstances, exploring and bringing them to new peaks, if necessary. In looser terms, one may say that the edge of chaos⁵ is the regime that provides an escape from the innovator’s dilemma (Christensen, 2013), at least in terms of balancing exploitation and exploration.

Edge of chaos regimes offer firms a better chance of reaching a higher performance than in rigid or chaotic regimes. Figure 4-2 depicts a qualitative representation of the performance levels that can be reached in different regimes and in different ‘areas’ in or near the edge of chaos. Performance tends to be higher at intermediate (internal) complexity levels, as work on the NK-model (Kauffman, 1993) has demonstrated. In summary, using the landscape metaphor, this effect results from the fact that with higher complexity, and therefore more finetuned component assemblies, average peak height is increasing. However, if complexity is increasing, the landscape becomes more rugged, making it harder for firms to oversee the landscape and identify the highest peaks – the more complex, the harder innovation becomes, the lower the average performance that can be achieved.

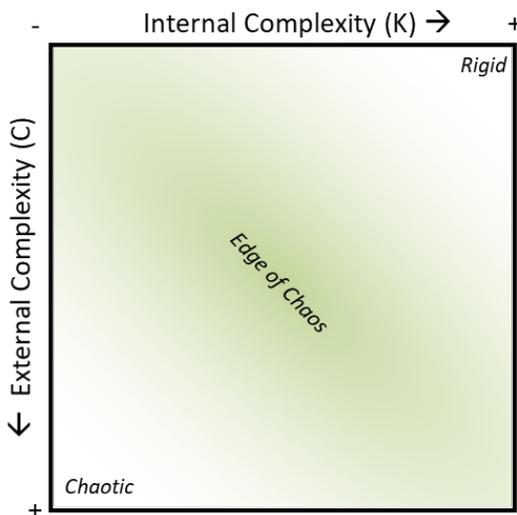


Figure 4-2. depiction of different ecosystem regimes in the area that is defined by internal and external complexity levels going from low to high, following table 4-1. The color intensity is used to represent an approximation of relative performance levels, with highest levels in the center of the picture; in the edge of chaos with intermediate internal complexity, balanced with corresponding external complexity.

5 Another way of understanding the edge of chaos regime is to consider the concept of ‘requisite variety’ (Ashby, 1991): internal complexity must be high enough to offer sufficient variety to deal with changes in the environment, as a result of external complexity.

Importantly, the original NKCS-model basically ignored the influence of the ecosystem network structure on the ecosystem regime. Instead, it simply assumed one specific network structure (a grid). However, as will be demonstrated below, the network structure is highly relevant, with clear differences between random, regular, small-world networks (Watts & Strogatz, 1998), and possibly other network structures, the latter being outside the scope of the current work. For alternative structures, see the work of Press (2008), in which predefined network structures were compared, notably clusters of inter-connected firms and ‘hub and spoke’ networks, in which firms are not connected to each other but to a leading or coordinating firm.

Short average firm-to-firm distance like in random networks is expected to support adaptability or flexibility, pushing the ecosystem in a more chaotic direction, as innovations in one firm will impact other firms in a smaller number of steps. More clustering like in regular networks is expected to support stability or rigidity, pushing the ecosystem in a more rigid direction, as innovations in one firm tend to remain longer within its cluster. Propagating the impact to other firms outside the cluster will take longer on average, because of higher average firm-to-firm distance compared to random networks. Small-world networks may offer both stability (clustering) and adaptability (short firm-to-firm distance), supporting the ecosystem in reaching or approaching the edge of chaos regime.

4.5 Model Implementation

Innovation (Fleming & Sorenson, 2001b) and co-innovation are considered here as an evolutionary process with discrete time steps, working on a set of – initially randomly generated – configurations, one configuration for each firm and with all configurations having similar length N and internal complexity K , with each firm having a randomly assigned matrix of interdependencies between components. Configurations are modelled here as binary strings, or bitstrings: each component is represented by either 0 or 1, without loss of generality. External dependencies C between firms that are directly connected to each other in the ecosystem network, are assigned randomly for each firm-to-firm connection. Based on internal and external dependencies K and C , each firm keeps track of performance value assignments to specific combinations of internal and external component states – ensuring consistency in performance values when the same combination is encountered more than once by the same firm.

In the current model, firms innovate by doing exploitation (Nelson & Winter, 1982), or local hill climbing, meaning that in each hill climbing step, a single component is changed (from 0 to 1 or from 1 to 0), provided that the performance of the resulting new configuration is better than the existing configuration of the firm. In our model we have applied an algorithm that is known as ‘greedy search’ (Goldberg, 1989), meaning that a firm in each hill climbing step will select 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. The hill climbing steps, one in each time-step, are repeated until no more improvements can be found. At that moment, a firm has reached a local (or the global) optimum – there is no way that exploitation can further improve a

firm's performance. Once reached, this status may remain the same for one or more time-steps, but will end when a change occurs in one or more components of the other firms to which the firm at hand is directly connected in the ecosystem network, and on which one or more of its components are dependent. As discussed before, such a change will deform the performance landscape, at which point new local hill climbing may be necessary to find a local peak, again⁶.

In the original NKCS work of Kaufmann and Johnsen (1991) and often in other NKCS work (Vidgen & Bull, 2011), actors are positioned in a regular grid, representing the ecosystem network, with each actor having 4 neighbors north, south, west and east to its positions (edge actors having 3 neighbors, corner actors having 2 neighbors), basically ignoring the possible impact of this or other structures on results. See figure 4-3 for an example.

In the model of the current work this network structure of actors – in this case firms – is variable, to study the effects of different network characteristics on co-innovation, as explained before. To do this, the model was implemented by using different network types, following the work of Watts & Strogatz (1998). Starting with a regular network, other networks were created by randomly rewiring connections between firms, creating small world networks and finally fully random networks. In short, regular networks have the key characteristic of high clustering (firms' neighbors are often also neighbors of each other) but long average firm-to-firm distance. Random networks have short average firm-to-firm distance but low clustering. Small-world networks have both high clustering and short average firm-to-firm distances. See chapter 3 for a further explanation.

Varying the network structure within a NKCS-model also reveals a parameter that was implicit in the original NKCS work: the number of direct connections D or 'degree' of one firm with other firms. In the original work based on regular grids, $D=4$ ($D=3$ for edge firms, $D=2$ for corner firms). In the current work, D is initially set to 4 for regular networks, but with network rewiring, D becomes variable.

Parameter settings that were used for simulations are listed in table 4-2.

To validate the model, results of simulations were cross checked with earlier work in two ways. First, simulations were done with $C=0$, basically creating the equivalent to the NK-model (Kauffman, 1993) with multiple but isolated firms – since there are no dependencies, the network structure does not play a role in that case. Second, the original NKCS work (Kauffman & Johnsen, 1991) was replicated, for which the network structure was replaced by a regular grid or lattice. In both the NK and NKCS case, results were in accordance with the original work. In addition, a sensitivity analysis was carried out to check for robustness of results when parameters are changed.

Based on the model and parameters as described above, we have carried out a series of simulations, using parameter settings that are indicated in table 4-2 as 'current work'. To prepare simulations, a set of 100 networks of S firms for each degree D and rewiring factor β combination was generated, with each simulation randomly selecting one of these networks. Rewiring factor $\beta = 0$ was used to generate regular

⁶ Interestingly, with greedy search, the process may come relatively more easily into recurrent cycles (with periodic attractors), as firms evolve more deterministically and may return repeatedly to the same genotype, causing a repeating response pattern in other firms.

networks (no rewiring), $\beta = 1$ was used to generate random networks (full random rewiring). Given the fact that regular simulations were carried out with $S=25$ and $D=4$, implying a total of 50 links in the network, $\beta = 0.1$ was used to generate small world networks with 10% or 5 links rewired at random.

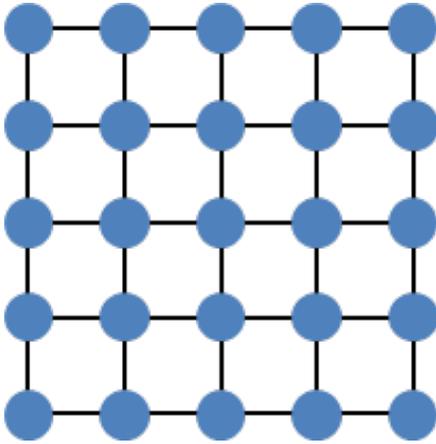


Figure 4-3. Example of a regular grid, in this case of 5 x 5 firms, with firms depicted as circles. Inner firms have 4 neighbors, firms on the edge have 3 neighbors, firms on corners only 2.

Parameter	Description	Settings NK equivalence (Kauffman, 1993)	Settings NKCS replication (Kauffman & Johnsen, 1991)	Settings current work	Settings sensitivity analysis
S	Number of firms in the ecosystem network	25	25	25	10, 25, 50
D	Degree (number of neighbors in regular network)	n/a	4 (3 for edge, 2 for corner firms)	4	4, 6, 8
β	Rewiring factor	n/a	n/a	0, 0.1, 1	0, 0.01, 0.05, 0.1, 0.2, 1
N	Number of components	16	16	16	8, 16, 24
K	Internal complexity	3, 7, 11, 15	3, 7, 11, 15	9, 11, 13, 15	For N=16: 0, 1, 3, 5, 7, 9, 11, 13, 15 For N=24: 4, 8, 10, 15, 22 For N=8: 1, 4, 7
C	External complexity	0	1, 2, 3, 4	1	0, 1, 2, 3, 4

Table 4-2. Parameters and their settings, used in simulations

Once prepared, simulations for all parameter combinations were carried out. In each simulation, firms were allowed to innovate in a series of time steps, limited to a maximum number of time-steps, after which no relevant changes in results were observable, as determined in preliminary experiments.

The model was implemented in NetLogo (Wilensky, 1999). For each simulation, with 100 runs for each parameter value combination, the model was executed according to the following pseudo code:

```

;initiation
initiate-network ( $\beta$ , S, D parameters)
initiate-NKCS-landscape (N, K, C parameters)
for firm f=1 to S
    bitstring (f) = random-bitstring (N=16)
    mutated(f) = true
next firm f
frozen = false
;run
for time-step = 1 to 500
    for firm f=1 to S
        mutated'(f) = mutated(f)
        mutated(f) = false
    next firm f
    for firm f=1 to S
        if one-of-direct-neighbour-firms-mutated'(f) = true then
            mutated(f) = greedy-hillclimb-successful(f)
        end if
    next firm f
    frozen = no-firm-has-mutated
next time-step

```

4.6 Results

Using the model as described above as a common baseline, several simulation experiments were carried out. First, the original NKCS-model was reproduced, using a regular 5x5 grid network of firms. In subsequent experiments, this network was replaced by regular, small-world and random network structures to study their effect on ecosystem regimes and associated firm performance levels. Network structures were also compared to see the effect of shocks being applied to ecosystems in frozen or equilibrium state, and to see under which regimes the resulting innovation cascades in their size and duration would follow a power law (Dosi, 1988). In addition, resilience characteristics of different ecosystem regimes and network structures were compared. Finally, experiments were carried out to see how network structures evolve over time, given the influence of these structures on the ecosystem regime and associated advantages and disadvantages in terms of performance and resilience.

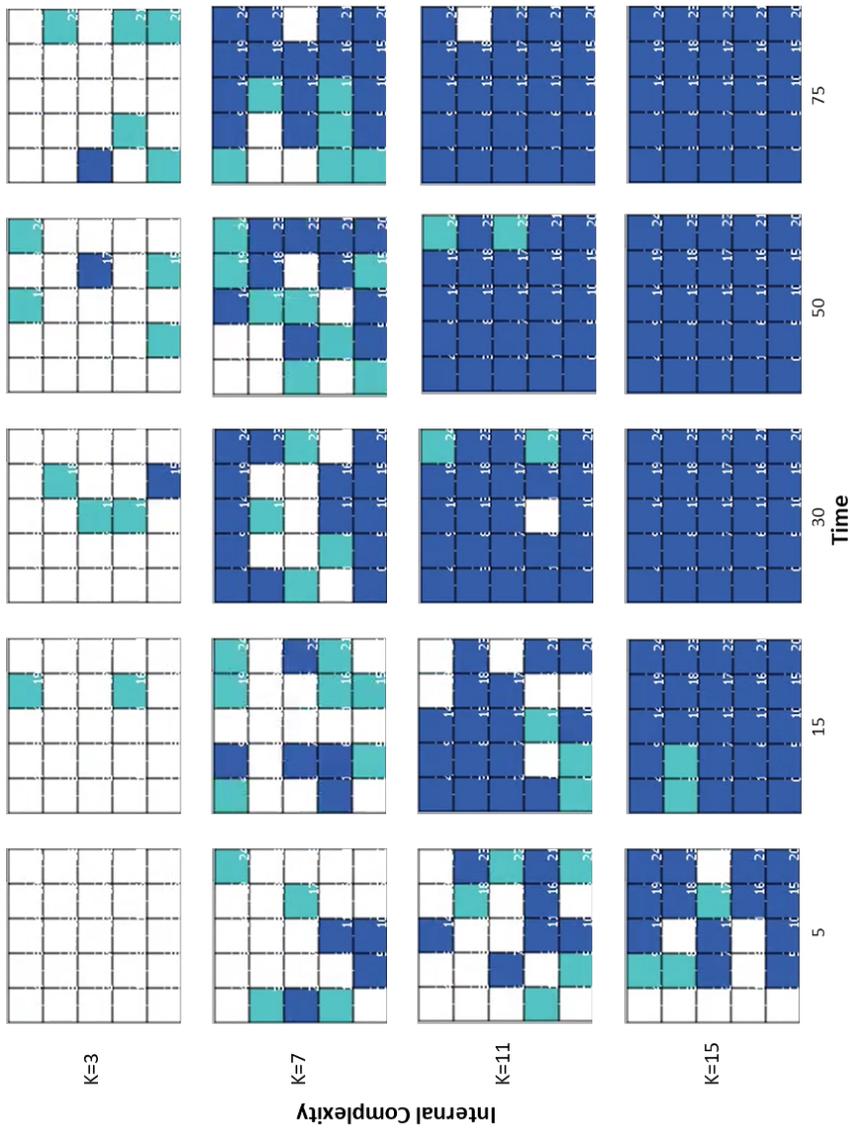


Figure 4-2. Example visualizations of NKCS simulations over time, in a regular lattice network structure (5x5 grid for $S=25$). In such a grid, most firms have 4 immediate neighbours, firms on the edge have 3 immediate neighbours and firms on a corner have 2 immediate neighbours. Firms are depicted here as a square, with colors white (firm innovated at that time step), light blue (firm did not innovate at that time step) or blue (firm did not innovate this and previous time step, or longer). In other words, the bluer the grid, the more the ecosystem is frozen or rigid, and vice versa. The examples shown here are for $N=16$, $K=3, 7, 11$ and 15 , $C=1$. The overall pattern is that for low internal complexity ($K=3$), firms keep on innovating, remaining in a 'red queen dance' (Van Valen, 1977) under a chaotic regime. For high internal complexity ($K=15$), firms quickly stop innovating, finding themselves in a rigid regime. For more intermediate internal complexity levels ($K=7, K=11$), firms keep on innovating longer, but still converge to a more stable or 'edge of chaos' regime, most so for $K=7$. See also results in figures 4-3 and 4-4.

4.6.1 *Reproducing NKCS grids*

Differences in ecosystem characteristics and performance levels can be explained by differences in internal and external complexity levels, as discussed in the previous section: high internal complexity (depending on K , relative to N) and low external complexity (depending on C , relative to N , and S , D and network structure) contribute to rigidity, high external and low internal complexity contribute to chaos, and *vice versa*. Reproducing the original NKCS-model in preliminary simulations on a regular grid of 5x5 firms confirmed these findings and assured consistency between the current and original model implementations. See figure 4-2 for example visualizations, demonstrating the different ecosystem regimes that may occur, in this case for different levels of internal complexity.

In fact, the results in the current study extend the evidence that was presented in the original work, and is therefore providing further underpinning of the seminal work by Kauffman & Johnsen (1991). See figures 4-3 and 4-4 for an illustration of results: higher C pushes the ecosystem towards more chaotic regimes, with firms continuing to innovate in response to changes in other firms, never reaching a Nash equilibrium. And *vice versa*: lower C decreases the chaotic nature of the regime, with $C=0$ corresponding to isolated firms that have no interdependencies with other firms, in effect doing their innovation in isolation. Likewise, higher K push the ecosystem towards more rigid (or less chaotic regimes). In the case of $K=3$ in figure 4-3, only chaotic regimes exist, with the exception of $C=0$. In case of $K=7$, it seems that with $C=1$, the ecosystem is closest to edge of chaos, balancing rigidity (allowing firms to converge towards a higher and stable performance level) with chaos (allowing firms to escape from lower performance levels at local optima). For $K=11$ and $K=15$, the regime for $C=1$ is becoming more rigid, although still offering better performance than in case $C=0$. For $K=15$, the regime for $C=2$ is becoming less chaotic, but not quite yet close enough to the edge of chaos to outperform more rigid regimes. In all cases, performance levels are highest when ecosystems are closer to the edge of chaos.

4.6.2 *Network structures and ecosystem regimes*

We now turn to the results obtained with alternative network structures. In particular, we look at regular, small-world and random networks, obtained by tuning the rewiring parameter from 0 to 0.1 to 1, respectively. The results confirm the relevance of the ecosystem network structure in the occurrence of different ecosystem regimes, as can be seen in figures 4-5 and 4-6⁷. In figure 4-5, the development in time of the number of changing firms is shown for different internal complexity levels (K), repeated for each of the network types. Rigidity is reached quickest in regular networks with highest internal complexity ($K=15$), with firms halting their innovation very early. Chaotic regimes are most prominent in random networks with lowest internal complexity ($K=9$), with firms ever responding to each other's changes. Regimes that are closest to the edge of chaos, with changes continuing rather long, but gradually decreasing and avoiding chaos, yield the highest performance levels as shown in figure 4-5:

⁷ Note that all results in figures 4-5 and 4-6 are with $C=1$. For higher C -values, not shown here, chaos increases and edge of chaos is only reached, if at all, for maximum K -values.

superior performance is reached with regular networks for $K=11$, with small-world networks for $K=13$ and with random networks for $K=15$. This shows that given a certain K -value, there exist an optimal network structure that finds the balance between stability and flexibility characteristic of the edge-of-chaos regime. This is in line with the work of Press (2008), who also found that the optimality of network structures, although only for two static variants, varies with K .

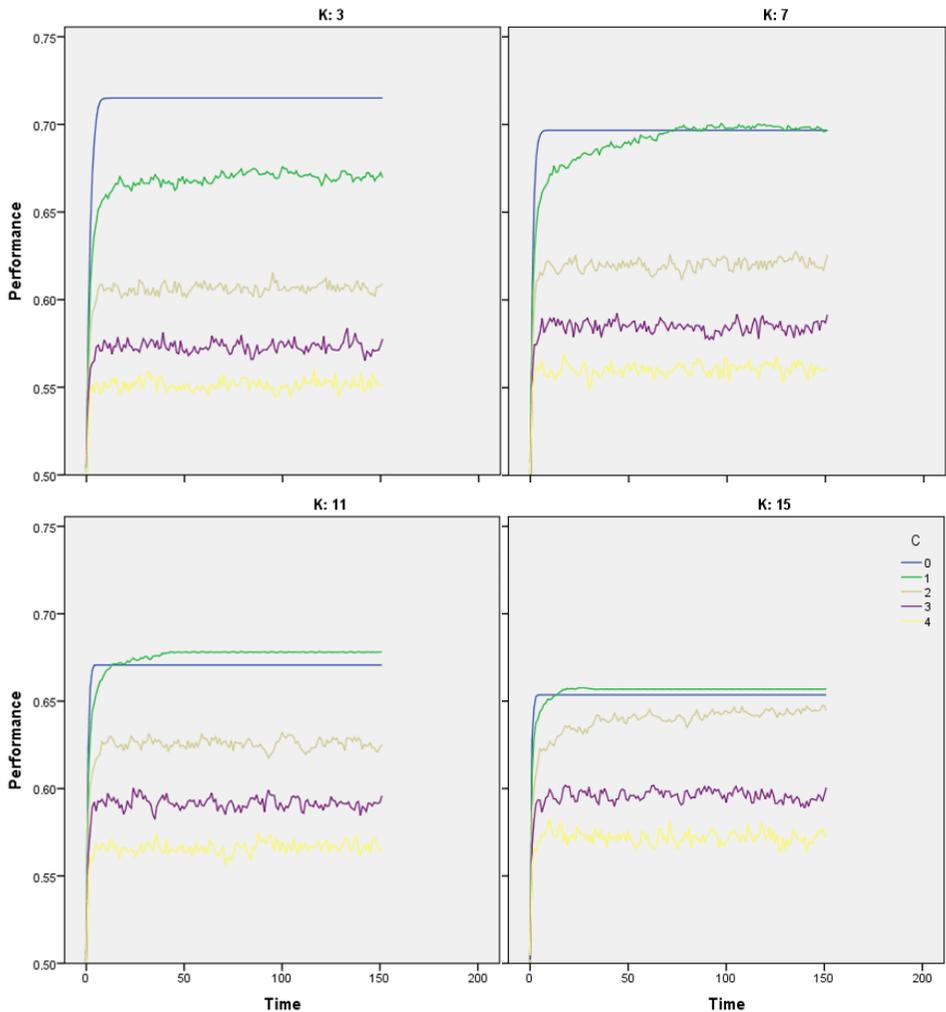


Figure 4-3. Performance over time, in a regular lattice network structure (5×5 grid for $S=25$) for $N=16$, $K=3, 7, 11$ and 15 , $C=0, 1, 2, 3$ and 4 . Performance levels are higher when the ecosystem is closer to edge of chaos regime, or lower when the ecosystem is increasingly too rigid or too chaotic.

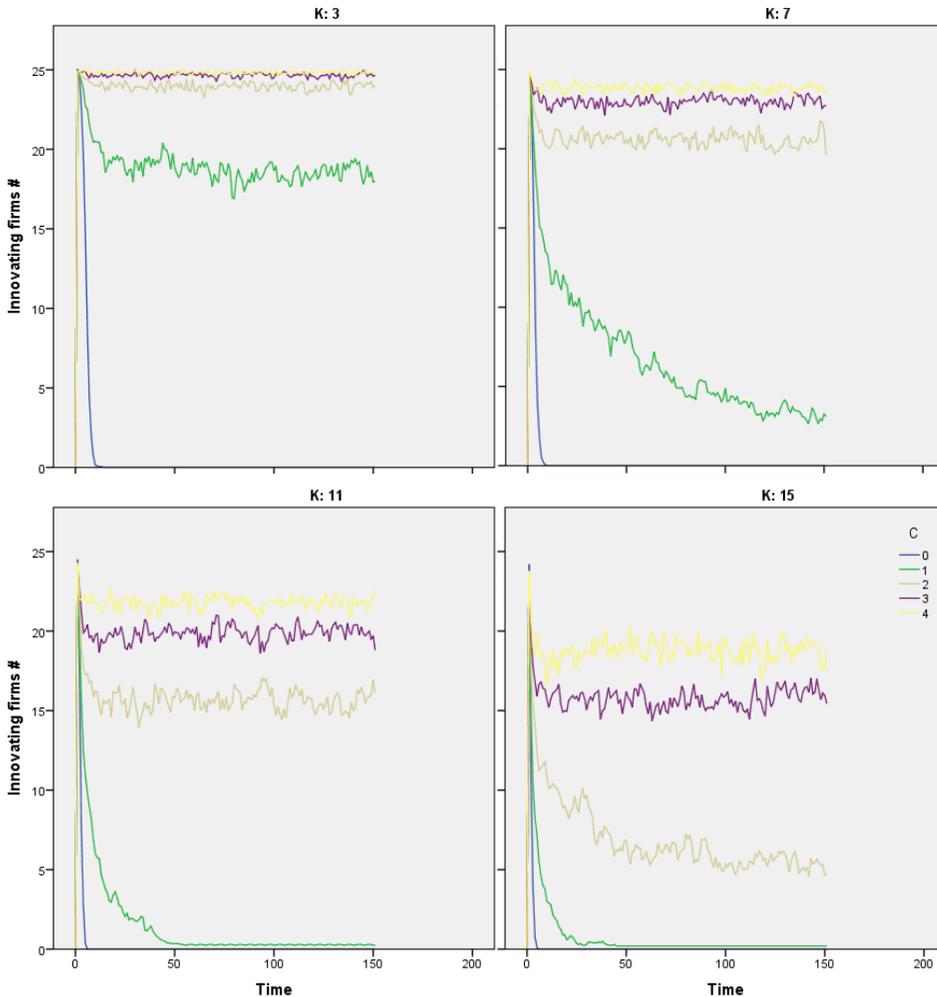


Figure 4-4. Number of innovating firms over time, in a regular lattice network structure (5x5 grid for $S=25$) for $N=16$, $K=3, 7, 11$ and 15 , $C=0, 1, 2, 3$ and 4 . In chaotic regimes, firms continue to innovate in response to changes in other firms, never reaching a Nash equilibrium. In rigid regimes, firms quickly stop innovating, as a result of high internal complexity (high K) or a lack of external dependencies (low C), both of which inhibit innovation. Near the edge-of-chaos, the number of innovating firms gradually decreases ($K=7, C=1$ or $K=15, C=2$), eventually coming to a halt ($K=11, C=1$), or ending in oscillations which result from repeating patterns of the same firms responding to each other's changes.

Considering these simulations, we observe that more random networks seem to invoke more chaotic regimes and *vice versa*: more clustered network seem to invoke more frozen regimes. In effect, network structures seem to ‘move’ the position of the edge of chaos. This can be understood when considering the fact that in more random networks, the average path length or distance between any two firms on

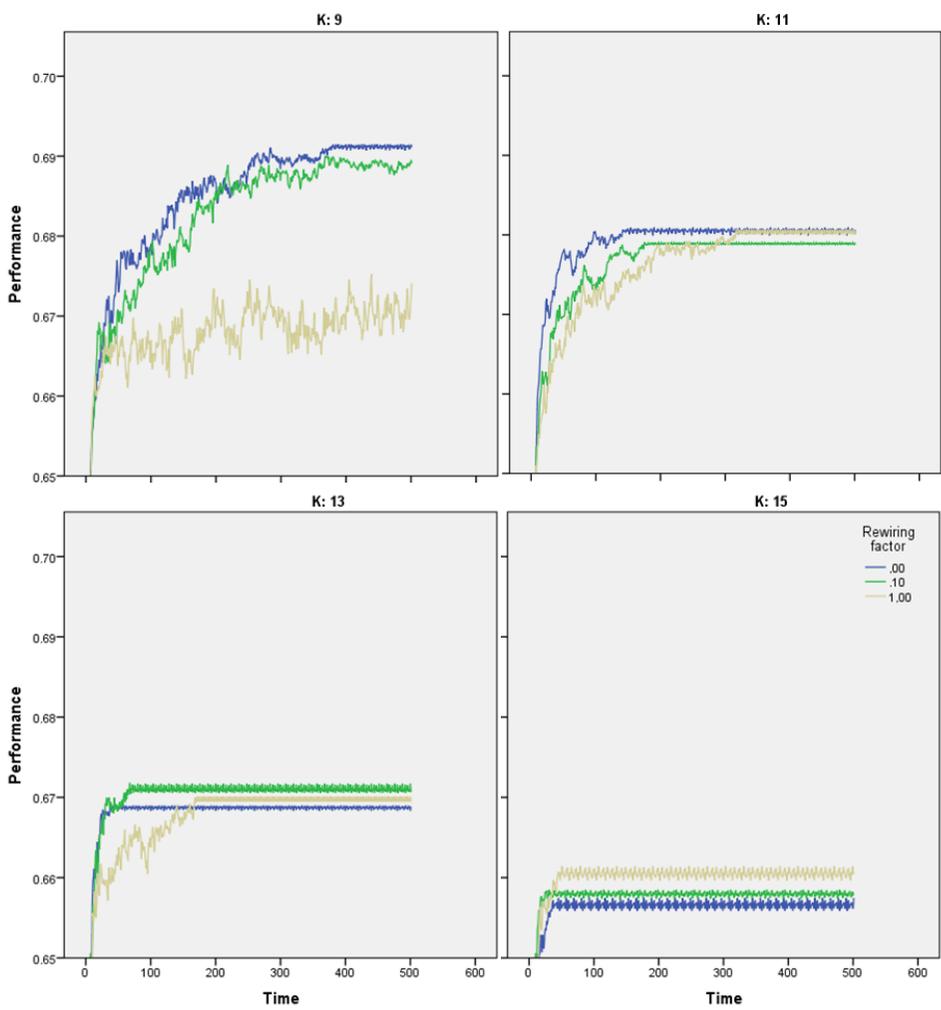


Figure 4-5. Rigid, edge of chaos and chaotic regimes for different levels of internal complexity (K) and different network types, all with N=16, S=25 and C=1, shown here as the change in performance (vertical axis) over time (horizontal axis). Some simulations end in oscillations which result from repeating patterns of the same firms responding to each other's changes, which is reflected in the averages shown here. Performance levels are highest for regimes that are closest to the edge of chaos.

average is shorter than in more regular networks. In more regular networks, firms are more clustered on average: firms are connected to other firms that on average are also more often connected to each other. These differences between network types are relevant when a firm innovates, impacting its direct neighbors through C dependencies between their components, and these neighbors impacting their neighbors, and so on. In other words, the impact of the original innovation ‘travels’ through the network, with a speed and duration that is not only influenced by

parameters like C and K, but also by the network structure. In a random network with shorter average network distances between firms, the impact will affect more firms in a smaller number of steps and therefore more quickly, thus pushing the ecosystem in a more chaotic direction, relative to more regular networks. And *vice versa*, in a more regular and clustered network, the impact will tend to remain more inside a cluster and will take longer to reach other more remote firms, thus pushing the ecosystem in a more rigid direction, relative to more random networks.

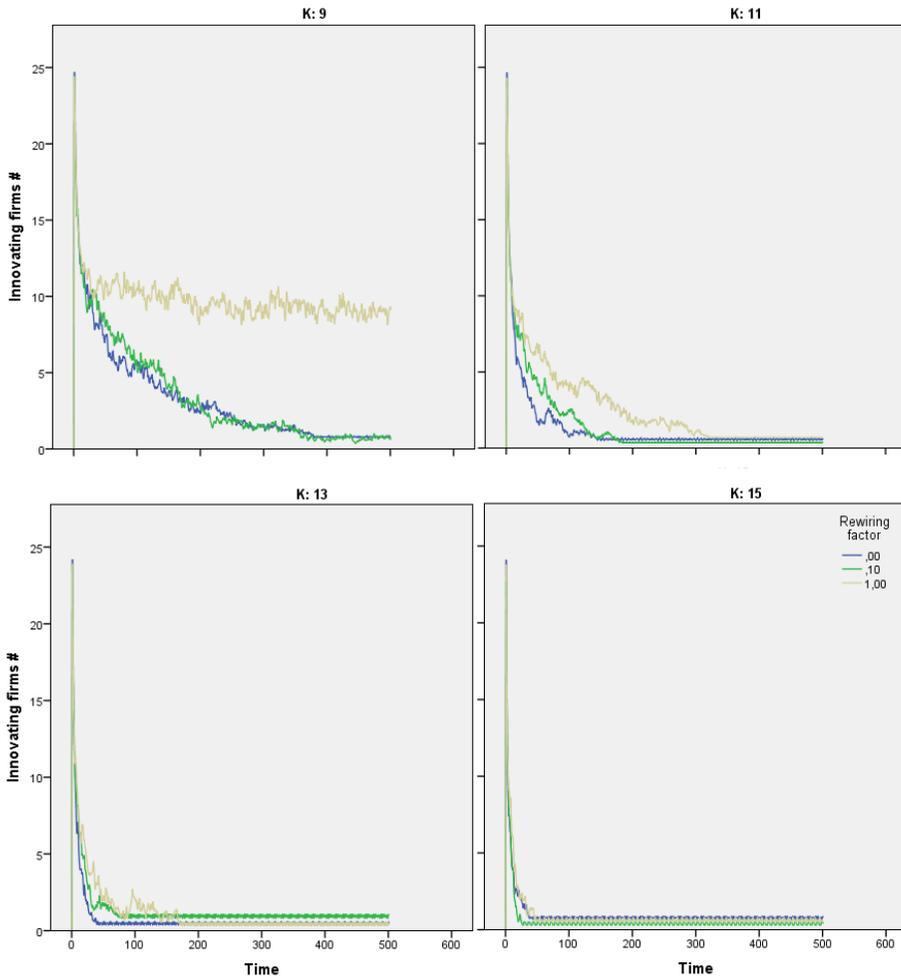


Figure 4-6. Rigid, edge of chaos and chaotic regimes for different levels of internal complexity (K) and different network types (rewiring factor 0 (regular), 0.1 (small-world) and 1 (random), all with N=16, S=25 and C=1, shown here as the change in number of innovating firms (vertical axis) over time (horizontal axis). On one extreme, K=9 and random networks have a chaotic regime. On the other extreme, K=15 and regular networks have a rigid regime. In between, for instance K=11 and small-world networks, seem to have edge of chaos regimes: the number of innovating firms declines over time, unlike chaotic regimes, but not so fast, unlike rigid regimes.

R square values						
	K=9	K=11	K=13	K=15		
Regular	0.7677	0.9647	0.9444	0.9422		Rigid
Small-world	0.4905	0.9484	0.9567	0.947		Edge of Chaos
Random		0.7098	0.9172	0.9738		Chaotic

Table 4-3. In figure 4-7, a linear trend line is visible for each group of observations of cascades, plotted on a logarithmic scale. R-square values indicate how close each line is to observations: the higher the R-square value, the closer observations follow a power law distribution. This table, with colors representing regime, provides an overview of R-square values, further clarifying that innovation cascade size distributions are close to power law distributions in edge of chaos regimes.

4.6.3 Network structures and power laws

To study the occurrence of innovation cascades or avalanches (Silverberg & Verspagen, 2005; Arthur, 2007; Dosi, 1988), possibly ranging from small scale innovation trains to large scale transitions, the model is extended by a mechanism to invoke such cascades in the ecosystem. Once ecosystems become ‘frozen’ – with Nash equilibria reached for all firms – a “shock” is applied by randomly reinitializing the configuration of the firm with lowest performance at that time. In real-world ecosystems, this may reflect the situation in which poor performing firms are more likely to do a radical innovation as an ultimate way to stay in business, or that such firms disappear and are replaced by new firms with a different approach. The applied shock causes all connected firms to respond, possibly causing yet other firms to respond, and so on. By measuring the time between shock and the moment the ecosystem becomes frozen again, and by measuring the number of innovations by firms during that interval (possibly with more than one innovation by a single firm), the size of the cascade is determined, expressed here as ‘time steps x number of innovations’. This process of freeze and shock is repeated for a number of times in each simulation, until a predetermined number of 1000 time steps has been reached.

Comparing different possible states of the ecosystem (figure 4-7), we observe that an increasingly chaotic state (for low K) will lead to more and ever larger cascades and fewer small cascades, with endless cascades once chaos prevents the ecosystem from ever reaching a frozen state. An increasingly rigid state (for higher K) will lead to more and ever smaller cascades and fewer larger cascades. Consequently, innovation cascade sizes in both chaotic and rigid states do not follow a power law distribution. It is only in edge-of-chaos regimes that the occurrence of smaller and larger cascades is close to a power law distribution and with a good statistical fit (Table 4-3) (West, 2017), which is consistent and further underpinning the original NKCS-model of Kauffman & Johnsen (1991) and the empirical distribution of innovation cascade sizes (Dosi, 1988).

In contrast, size distributions in more chaotic regimes tend to be skewed to larger sizes, and in more rigid regimes to smaller sizes, with the regime not only determined by K and C, but also by the structure of the ecosystem network. With smaller cascades happening relatively more often in clustered ecosystem networks, and bigger cascades happening relatively more often in random ecosystem networks, and vice

versa: clustered networks seem to inhibit bigger innovation cascades and transitions, random networks prevent innovation cascades from staying small.

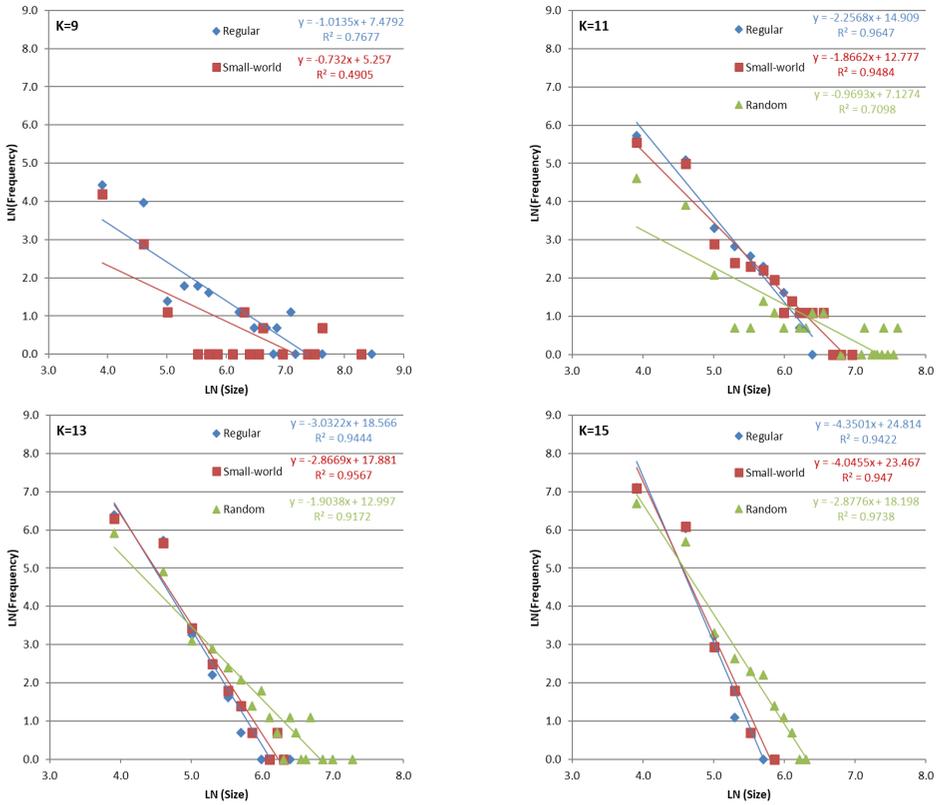


Figure 4-7. Distributions of innovation cascade sizes, logarithmic scales, for different internal complexity K levels. Observations are binned (size 100) for visualization purposes. Linear trendlines reveal that with regular networks & K=11, small world networks & K=13 and random networks & K=15, distributions are most similar to a power law distribution (highest R square values), corresponding to edge of chaos regimes and highest performance levels. In all cases N=16, C=1 and S=25.

The above results were obtained in simulations in which shocks were only applied once a frozen ecosystem state was reached. Alternatively, in other simulations, shocks were applied at certain time intervals, regardless of the state of the ecosystem, which in the real world may correspond to exogenous events that happen periodically, like an important scientific discovery that gives rise to a radical innovation, a change in climate or an economic hype or bust, creating a significant disturbance in the ecosystem. Results of these simulations (not shown here) were similar to what is reported above, providing further trust in our observations. An interesting difference, however, was that with more frequent shock application, the occurrence of edge of chaos regimes required more rigidity in the ecosystem, with higher K, lower C, or more regular networks. This seems to make sense if we consider that each shock creates a

cascade of change, small or big, thus pushing the ecosystem more towards a chaotic regime. In other words, following this line of reasoning, in an economically or otherwise volatile context, ecosystems that are more rigid (but not too rigid) are expected to perform better.

Interestingly, results also demonstrated that a moderate amount of shock application may be beneficial to improve performance levels, even in mildly chaotic regimes. This can be explained by considering the fact that each and every shock introduces some 'newness' in the ecosystem, invoking the exploration of other parts of the search space for one or more firms. However, as mentioned above, too frequent shock application will push the ecosystem towards more chaotic regimes, ultimately leading to destabilization of the search space, in which booking sustained progress becomes impossible.

The occurrence of (near) power law distributions of innovation cascade sizes, only seems to happen in regimes that are at the edge of chaos or close to it, associated with highest performance levels. In other words, the existence of a (near) power law distribution seems to reflect the underlying 'health' or 'quality' of the underlying ecosystem (West, 2017), balancing stability with adaptiveness, allowing actors to reach their highest performance levels. In other words, innovation cascades that follow a power law distribution in their size, are a characteristic of ecosystems that are on the edge of chaos, with the right mix of internal and external complexity (K and C), number of actors (S), degree (D) and an appropriate network structure (according to rewiring factor β).

4.6.4 Resilience

Firms benefit from reaching higher performance levels when their ecosystem is in or close to edge of chaos regimes. However, ecosystem 'health' is not only defined by reaching performance levels that are as high as possible (Roundy et al., 2017; Iansiti & Levien, 2004). To sustain higher performance levels, ecosystems must also be able to deal with disturbances and shocks, which are likely to have a negative impact, small or large, on performance. In other words, ecosystems need to be resilient in order to limit the effects that shocks may have (Roundy et al., 2017). To do so, ecosystems may have characteristics that minimize the initial impact of shocks, or that help overcome the impact of the shock or 'repair the damage' as soon as possible, or both.

Here, we measure resilience by looking at the extent to which ecosystems are able to restore or perhaps even improve performance, after a shock has occurred, and how much time that takes. This is done by measuring the average performance level of all firms in the ecosystem, just before a shock occurs: the pre-shock ecosystem performance level. Then, after a shock occurs – leading to a small or large cascade of change and, in general, initially lower performance levels – the average performance level is monitored until it is at least back at the pre-shock level, at which point the duration to restore performance can be established.

In the results that are presented below, shocks were applied to randomly selected firms by randomly reinitializing the configuration. In terms of the NKCS model, this is equivalent to redrawing a random bitstring for each firm that was selected to be part of the initial shock. In turn, other firms that are directly connected to the firms in the initial shock, will experience a performance change and in response need to change their own configuration, which may affect still other firms, and so on.

In addition, to check for robustness, simulation experiments were carried out by applying different ways to select firms for the initial shock. Instead of just randomly selecting firms, shocks were applied to the poorest or best performing firms. These variants differed in terms of the impact the shock had on overall performance degradation. Nevertheless, we found that these differences created no qualitative differences in the simulation results.

Typically, the ecosystem and firms therein will need a certain period of time in which performance is gradually restored, in response to the initial shock. If the initial shock has a larger impact, then more time will be needed to restore performance, on average, and *vice versa*. This implies that to measure resilience, as mentioned, the duration to restore must be related to the shock impact size, which can be determined in several ways, taking different factors into account. Such factors include the initial change of performance level, caused by the shock, the number of firms that was involved in the following cascade, or the total time duration of the cascade.

In table 4-4 results are shown of the 'relative restore duration', which is defined here as the duration to restore, divided by the number of firms that was involved in the cascade that was triggered by the initial shock. In one extreme, the relative restore duration is longest when it takes very long to restore performance in case of small cascades. In the other extreme, the relative restore duration is shortest when performance is restored quickly, even in the case of large cascades. Similar results were obtained, with consistent but varying levels of differences between regimes, when different factors are taken into account to determine the impact size of the shock and subsequent cascade.

Figure 4-8 demonstrates the significant differences in resilience between ecosystem regimes. More chaotic regimes, for example with $K=9$, do not only suffer from a bigger loss of performance as a result of shock, they also have more difficulty or do not succeed at all in restoring performance. More rigid regimes, on the other hand, for example with $K=15$, lose less of their performance when shocks are applied, but it takes relatively long to restore their performance. Regimes that are closest to the edge of chaos, like with $K=13$ and a small world network structure, offer the best resilience: not only do they need less time to restore performance, they even benefit from shock application, resulting in higher performance levels than before.

As can be seen in table 4-4, the relative restore duration is longest, on average, in most chaotic regimes (random networks and $K=9$). This seems to make sense as improving or restoring performance levels is very difficult if not impossible when firms have to innovate on the heaving deck of their continuously changing performance landscape. However, the relative restore duration is not shortest, on average, in most rigid regimes (regular networks and $K=15$). At first sight, this seems to be counterintuitive as rigid regimes are expected to limit the impact of shocks, which indeed is reflected in the fact that cascade sizes tend to be smaller in rigid regimes, as demonstrated in the former section. However, not only do we relate the restore duration to cascade size, we also have to take into account that regimes that are too rigid, make it harder for firms to find a way to improve their performance. Hence the adjective 'rigid': in very rigid regimes it will take longer to restore fitness. Considering this, it seems again that edge-of-chaos regimes offer optimal conditions, which indeed is confirmed in the current results, as demonstrated in table 4-4, in which

lowest relative restore duration (=best resilience) correspond with (close to) edge of chaos regimes (as indicated by colors). From this it can be concluded that edge of chaos regimes, with (near) power law distributions of cascade sizes, not only offer best performance levels, but also best resilience in the current interpretation of relative duration to restore performance levels after shock occurrence.

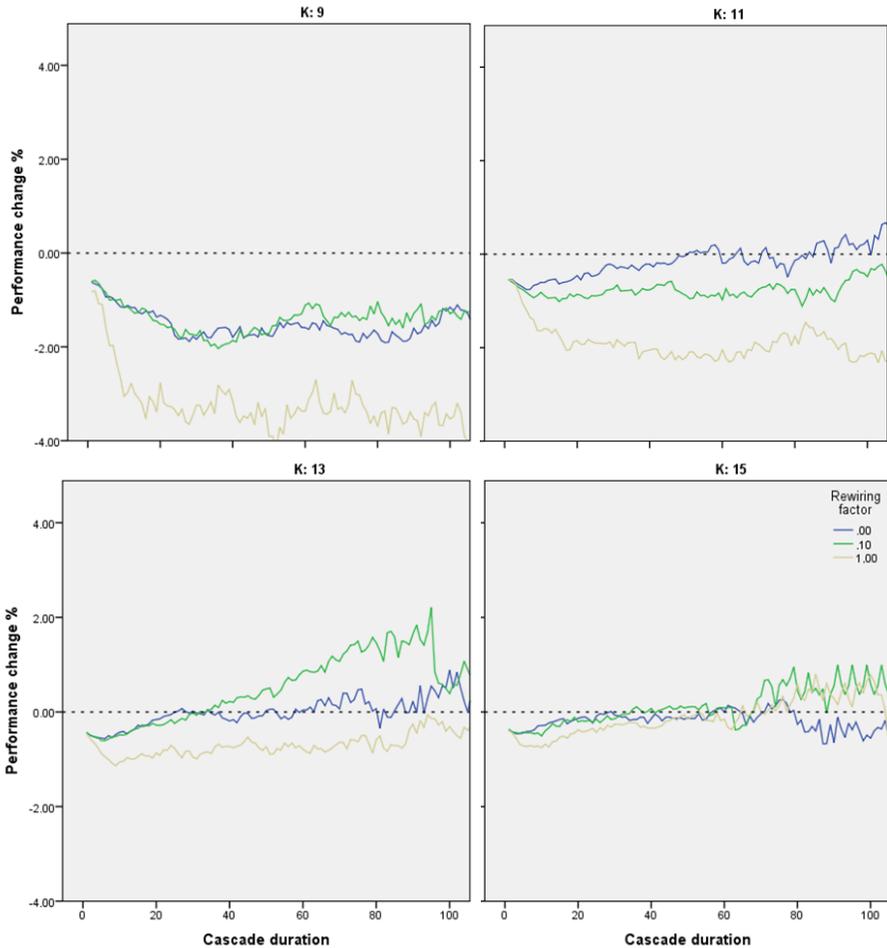


Figure 4-8. Resilience, represented here as the change of performance after a shock has been applied to the ecosystem and during the resulting cascade of innovations, for different network structures and different internal complexity levels (K), in all cases with $N=16$, $S=25$ and $C=1$. Results are averaged over all occurring cascades. As longer cascades occur less often than shorter cascades (depending on the ecosystem regime – see former section), the number of experimental observations decreases with cascade duration, which is therefore limited here to 100 time steps. Resilience is best in regimes that are closest to edge-of-chaos, both in terms of time and performance restoration or improvement.

Relative Restore Duration								
	K=9	K=11	K=13	K=15				
Regular	1.1	0.98	0.97	0.99		Rigid		
Small-world	1.34	0.95	0.9	0.94		Edge of Chaos		
Random	2.13	1.18	0.91	0.88		Chaotic		

Table 4-4. Resilience measured as relative restore duration (duration corrected for the initial shock impact size). Colors represent regime. Resilience is best when relative restore duration is lowest, which is the case in (close to) edge of chaos regimes.

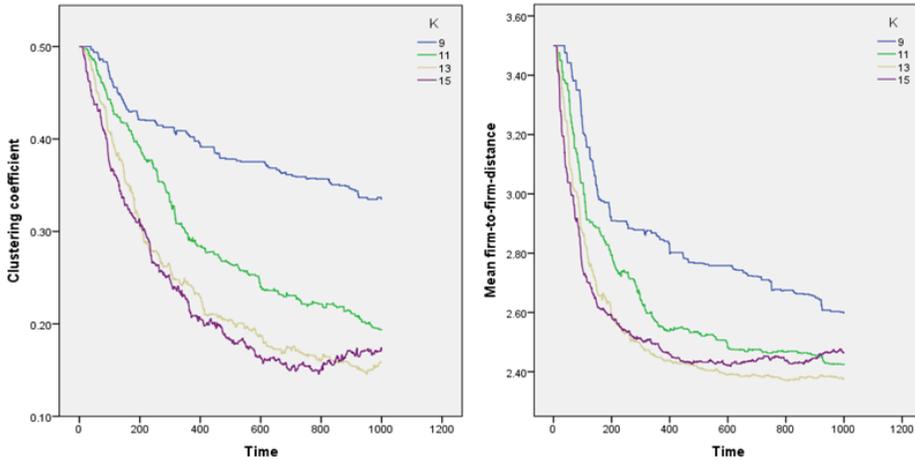


Figure 4-9. Evolution over time of clustering and average firm-to-firm distance in an ecosystem network with N=16, C=1 and S=25, with new firms applying ‘greedy linking’. For lower K values, clustering and firm-to-firm distances remain higher for longer, supposedly to maintain rigidity or to avoid chaos, in which performance levels are lower. Vice versa, for higher K values, less clustering and shorter firm-to-firm distances are desirable, escaping from too much rigidity.

4.6.5 Evolving ecosystem network structures

So far, we have assumed that ecosystems and characteristics like their network structure are a given. However, in reality both biological and business ecosystems evolve over time. Ecosystems grow or shrink in terms of the number of actors S, and with actors evolving themselves, changes⁸ occur in K, C and in the network structure, as actors may also change their interdependencies with other actors.

From an evolutionary perspective, both higher performance and better resilience could imply that ecosystems, including their K, C, S, D and β characteristics, would evolve towards the edge of chaos. However, the evolution of the ecosystem as a whole should emerge from the evolution that takes place at the level of individual actors,

⁸ It should be noted that the current NKCS-model would need to be changed to accommodate actor specific K and C values.

as there is no selective pressure at ecosystem level and ecosystems in general do not have supervisors or coordinators that can dictate these characteristics, although innovation policy makers and large dominant firms may have some influence. In other words, if we assume that actors are able to evolve or make changes to their internal and external complexity and the way connections are created with other actors, we may also assume that the ecosystem as a whole will evolve, reflecting the changes made at the level of individual actors.

Here our primary interest lies with how ecosystem network structures would evolve over time (Jain & Krishna, 2001). To further our understanding, the model was altered to simulate the evolution of a network of co-innovating firms. Starting with a regular network of S firms – other variants⁹ were also tested, with similar results – firms are allowed to evolve like in the former model, until a frozen state is reached. Once this is the case, the worst performing firm is removed, including its network connections with other firms, and a new firm is added to the network. This new, randomly initiated firm is connected to other existing firms, using what is called here ‘greedy linking’. In greedy linking a new firm selects existing firms for connection that have the most positive effect on the performance of the new firm. In non-greedy linking, new firms connect at random with existing firms. Preliminary simulations showed that greedy linking, compared to non-greedy linking, resulted in reaching a stable network structure more quickly, but with similar end results.

The process of evolving until a frozen state was reached and then destroying one firm and introducing a new firm, was repeated in simulations for 1000 time steps. As frozen states are more quickly reached with a more rigid ecosystem as a result of higher K or lower C , we expected to see faster ecosystem evolution in the case of this model when K is high, which is indeed observed in our results. Figure 4-9 demonstrates how, in this model, ecosystem network structure evolves over time, in this case showing the average local clustering coefficient¹⁰ (Newman, 1999) or ‘clustering’ for short, and the average firm-to-firm distance in the network. Not only do we see that for higher K , these characteristics start to change sooner and more rapidly, but also – importantly – level out at a lower level. Apparently, with higher K and hence more rigidity, the network must be less clustered (with shorter average firm-to-firm distance) to move towards the edge of chaos and reach higher performance levels. Put differently: the more rigidity, the more network randomness (with low clustering and shorter distances) is needed to escape from freezing, or to reach the requisite variety (Ashby, 1991).

These results indicate that ecosystem network structures may indeed evolve over time, moving towards the edge of chaos, emerging from the adaptation and selective pressure at the level of individual actors. Other network and ecosystem characteristics may emerge and evolve similarly (Jain & Krishna, 2001), like the ratio between internal and external complexity levels K and C , which however, we leave to future work.

9 One variant was to start with a random network. Another variant was to start with a small network of 3 firms, removing the worst firm once in a frozen state and adding two new firms. Thus allowing the ecosystem to grow to a maximum of S firms.

10 Indicating the extent to which firm neighbors are also neighbors of each other. Local clustering is highest in regular networks and lowest in random networks. See chapter 1 ‘introduction’ for a further explanation.

4.7 Discussion

The model that we have applied in this study, being in essence a combination of the seminal work of both Watts & Strogatz's work on networks (1998) and Kauffman & Johnsen's work on coevolution (1991), using just a limited set of parameters, seems to be able to reproduce and provide coherence to multiple –until now seemingly unrelated– stylized facts (Silverberg & Verspagen, 2005; Dosi, 1988) and empirical findings regarding small-world network structures between firms (Ozman, 2009) in existing work, strengthening our confidence in the results of our simulations and the answers they provide to our research questions.

It has proven useful to take the perspective of the economy as an ever changing web of interacting actors in which new or improved technologies enter continuously, triggering changes in other technologies or driving others extinct. More specifically, considering business ecosystems as complex adaptive systems with co-innovation as an co-evolutionary process (Fleming & Sorenson, 2001b), has provided more quantifiable insights about time-clustered innovations and the occurrence of power law distributions in the size of innovation cascades (Dosi, 1988).

Past work (Kauffman & MacReady, 1995) already provided the insight that business ecosystems may operate under quite different regimes. A rigid regime occurs when firms get stuck because of high internal complexity or because of a lack of external changes that would make change necessary. Chaos occurs when firms find themselves on a heaving landscape of ever changing performance levels, as a result of heavy interdependence between firms, on which making sustainable innovation is infeasible. Firms and the ecosystem at large operate best in a regime at the edge of chaos, comparable to a state of self-organized criticality (Bak et al., 1988), in which we imagine the system to be a pile of sand, with individual sand grains dropping on top: most grains lead to small changes or minor avalanches, but some may lead to massive avalanches, the size of these avalanches following a power law distribution (West, 2017). In the edge of chaos, stability and adaptiveness are balanced, providing the requisite variety (Ashby, 1991), but not too much.

The current work contributes by providing new insights about the important role of ecosystem network structure to avoid both rigidity and chaos, or to get nearer to the edge of chaos regime. Next to the already known relevance of internal dependencies (K), external dependencies (C), the number of actors (S) and their network degree (D), we found that the extent of clustering and firm-to-firm distance were significant in establishing a certain ecosystem regime. Highly clustered or regular networks seem to be most useful in getting closer to the edge of chaos when other characteristics tend more strongly to a chaotic regime. *Vice versa*, fully random networks without clustering and short average firm-to-firm distances are most useful when other characteristics pull the ecosystem towards a rigid state. Small-world networks, combining clustering with short average firm-to-firm distances, seem to be most useful in all other, less extreme cases, where the network structure helps to bring the ecosystem closer to the edge of chaos. Following this reasoning, it seems to make sense that in many empirical studies on networks on firms, small world networks are indeed observed (Baum et al., 2003; Chen & Guan, 2010; Ozman, 2009; Schilling & Phelps, 2007; Ter Wal, 2013), although not in all cases, possibly when more random or

more regular networks are needed to avoid rigidity or chaos, respectively. In addition, to explain empirical findings it should also be noted that networks of firms also exist for other reasons than co-innovation. Such reasons include knowledge sharing – see chapter 2.

Results from the current study further confirm earlier findings (Kauffman & Johnsen, 1991) that innovation performance levels are maximized at the 'edge of chaos': firms do not get stuck on suboptimal performance levels, nor does chaos degrade their performance too much. Importantly, adding to existing work, it was established that edge of chaos regimes also provide best resilience. Resilience that is needed when inevitably ecosystems have to endure shocks that result in performance degradation. To have sustainable performance levels, such degradations must be limited and repaired as soon as possible, which in edge of chaos regimes is most optimal.

4.7.1 Implications

In the real world, the optimal edge of chaos may be achieved by having some kind of coordination between firms or through innovation policy measures. But given the open and complex nature of ecosystems, making top-down coordination and policy effectiveness very hard, it seems more likely that this state must be reached by a process of self-organized criticality (Bak et al., 1988), emerging from the behavior of individual firms that strive to innovation performance maximization by implicitly or explicitly managing their individual internal and external complexity and their position in their business ecosystem network structure, as far as possible. An example of such self-organization seems to be found in an empirical lifecycle study in the microcomputer industry (Marion, 1999), using a NKCS-model perspective. This revealed that this sector was initially chaotic, with many small players (high *S*), but then became more stable, yet never rigid, with fewer and larger players and increasing competition (higher *C* and more network connections), due to globalization and other factors. Initially, internal complexity *K* was low, but started increasing and finally became lower again, due to standardization and modularization (Baldwin & Clark, 2000; Fleming & Sorenson, 2001a; Henderson & Clark, 1990).

The results of the current work may also be used to progress a general awareness amongst management scientist and practitioners that ecosystem characteristics have a significant impact on (innovation) performance of individual firms. If not literally by changing *K*, *C* and network structures, innovation managers and policy managers may benefit from understanding, monitoring and influencing business ecosystems in a more qualitative way. For instance, balancing stability and adaptability may also be supported by other means, like managing innovation to minimize or maximize responses to external changes explicitly, or like initiating new innovations, ignoring internal or external dependencies, as far as possible. After all, ecosystem actors like firms are not like biological organisms in the sense that firms and their people have a free will that allows them to deviate from co-evolutionary patterns, provided they are conscious of the process. Still, most firms cannot afford for too long to ignore changes in their business ecosystem, such as quality changes in supplied products, new customer demands or competitor innovations, all impacting the performance and competitive strength of any firm.

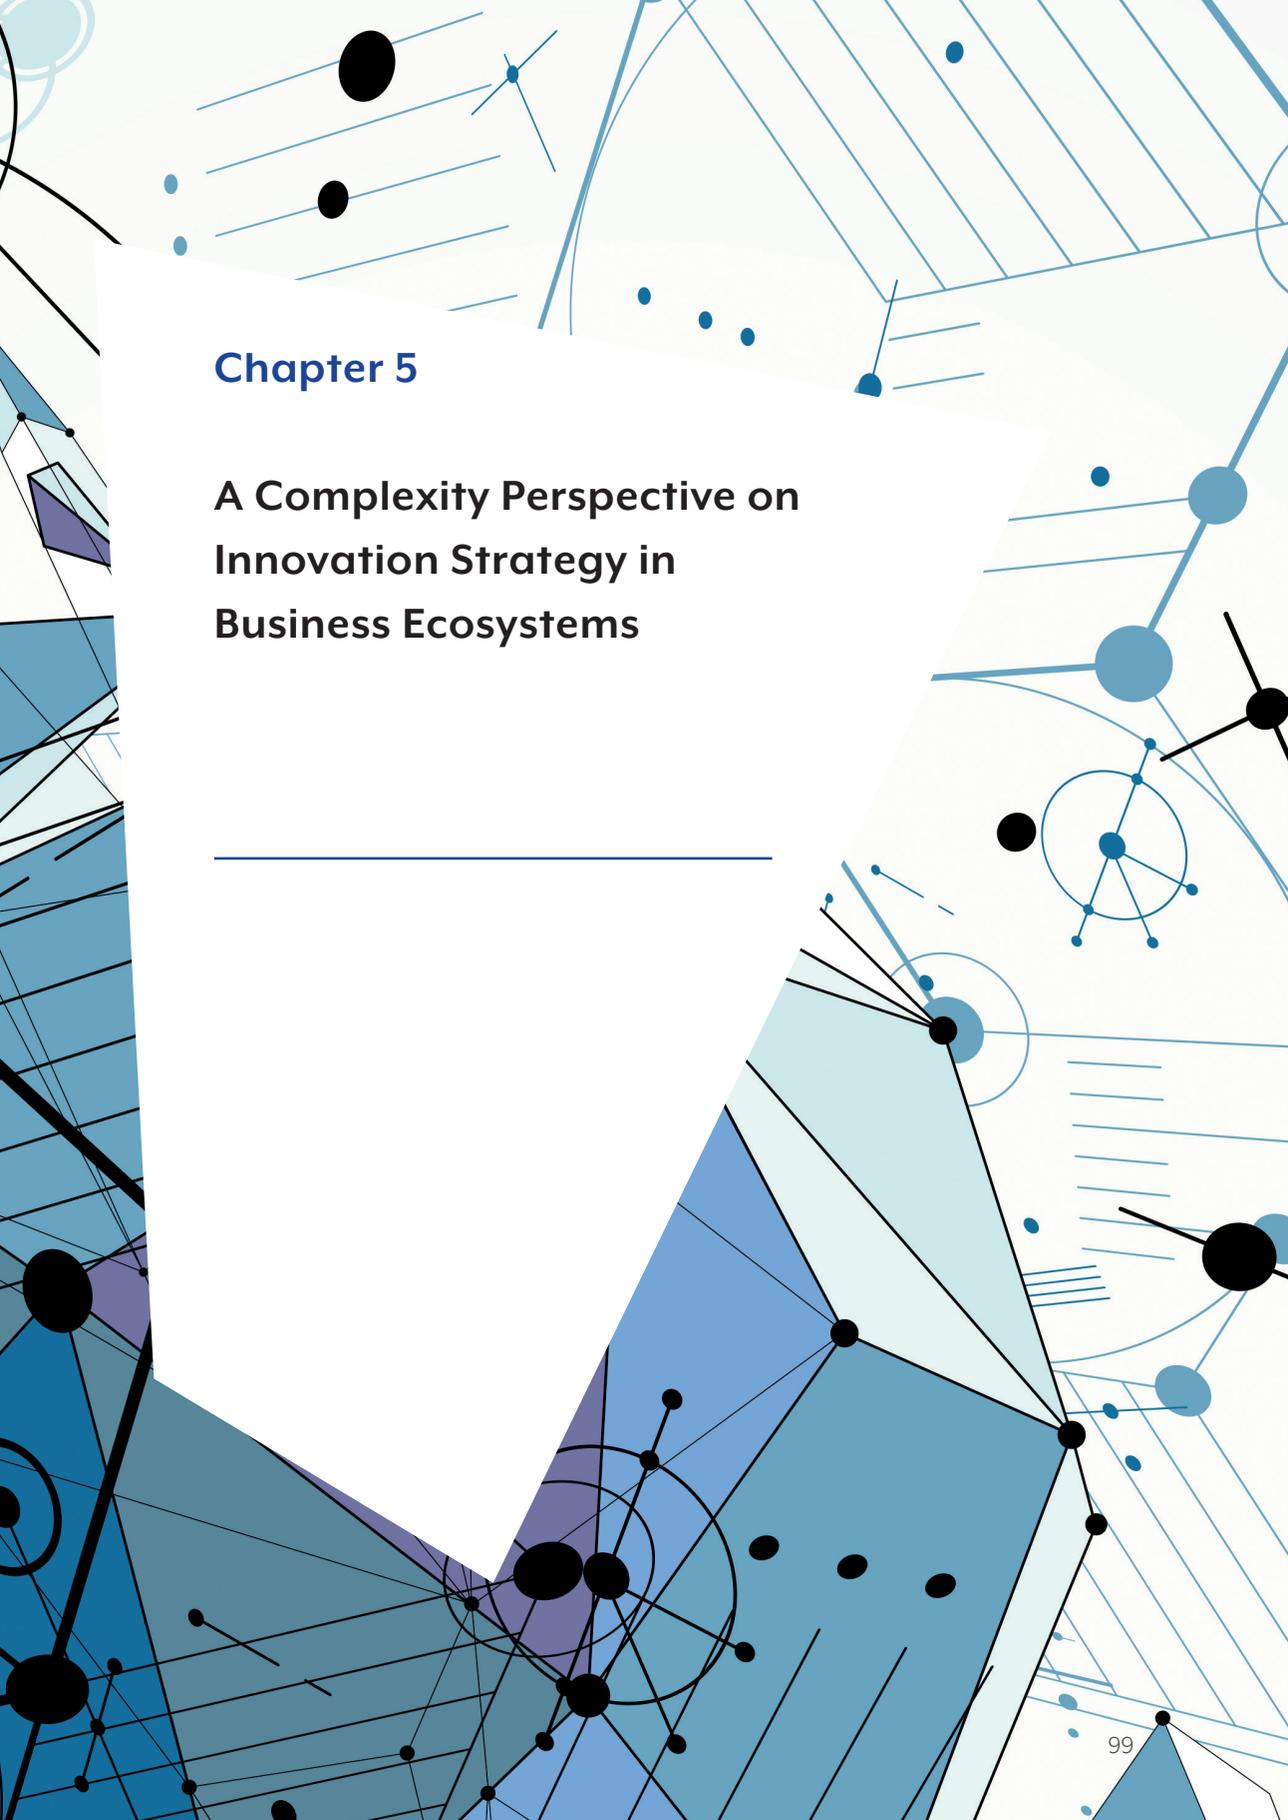
Firms may also use a more explicit awareness of interdependencies with other competing or complementary firms in their business ecosystem to coordinate and prioritize certain innovations. Traditionally, firms manage innovations primarily by considering the effect on their own performance, valuing independence, making the invalid assumption that firms can mostly do without others to do innovation (Adner, 2012). Using the perspective of co-innovation in business ecosystems, firms may also consider the effect of their innovation on the performance of other firms, and the other way around. Thus, complementary firms may benefit from coordinating their innovations, knowing that innovation in one firm has an impact on other firms. New innovations and corresponding innovation responses may be planned in time together, anticipating available resources, market conditions and other factors to increase efficiency and control on how and when interdependent innovations occur. If suppliers announce the nature and timing of their innovations with clients more timely, or engage in co-creation with their clients, these clients may prepare themselves in a better way by taking measures to manage their part of the innovation cascade, should they accept the innovation (or else switch to another supplier that continues to provide a component like it was before the innovation). Similarly, knowledge that is explicitly collected and maintained about interdependencies with competitors and about their innovation activities may be used to anticipate the necessity of new innovations. Such knowledge may also be used – more speculatively and in a less friendly manner – to orchestrate, possibly together with other firms, an innovation cascade that would create an unexpected heavy burden on competitors or even would create a chaotic regime, provided that such a regime could be kept contained to affect only competitors, deeming such a strategy as difficult and precarious. Nevertheless, knowing and understanding one's business ecosystem should be an increasingly important tool for firms to improve their innovation strategy and increase their competitive strength.

From a more centralized or governance policy making point of view, one may use the ecosystem co-innovation perspective in the interest of economic progress and stability. From this perspective, maximizing innovation performance in edge of chaos regimes may not be the only concern. For example, policy makers might also be interested in ecosystem resilience: indicating to what extent a business ecosystem is able to deal with shocks like firms that go bankrupt, causing a cascade of other firms falling over, similar to what happened in the global financial crisis that started in 2008 (Tedeschi et al., 2012). Or, more in the domain of innovation, shocks like a sudden change in consumer preferences or a strong increase of competition with new firms entering a local ecosystem as a result of globalization or digitization. As results from the current study demonstrate, if shocks are expected to occur frequently, then being on the rigid side of the edge of chaos seems to be the optimal regime, anticipating the chaotic push effect of shocks. However, we must limit our expectations about the real-world usefulness of such notions, as policy makers, let alone individual firms, are probably not able to gauge or orchestrate the regime of ecosystems very precisely.

4.7.2 Future work

The abovementioned example of the microcomputer industry (Marion, 1999) underlines the fact that ecosystems are not static. Firms come and go, the number of firms (S) growing or declining. Firms, alone or together, may adopt standards to decrease internal or external complexity, implying that the current model may also be used on studying standard-setting coalitions of firms (Axelrod et al., 1995). Firms may change the design or architecture of their products, resulting in more componentization (N) or modularization (Simon, 1996; Henderson & Clark, 1990), reducing complexity. And, as demonstrated in this study, firms may also consciously monitor their position in their network, and may change the way other firms are selected to work together, like selecting specific suppliers of components, if possible at all (Adner, 2012; Iansiti & Levien, 2004). Either way, changes at firm level may lead to other emergent properties at ecosystem level. Assuming that selective pressure is based on performance levels, it seems likely that ecosystems will evolve towards a state in which performance levels are maximized; the edge of chaos. However, developing a further understanding of how changes at firm level work out on ecosystem level, is recommended for future work.





Chapter 5

A Complexity Perspective on Innovation Strategy in Business Ecosystems

Abstract

Increasingly, the effectiveness of a firm's innovation depends on understanding and influencing its business ecosystem position, for which the ecosystem complexity model is introduced. Using internal and external complexity as dimensions, four ecosystem position categories are identified: completer, firms largely independent in their innovation, composers, firms offering a platform or infrastructure used by complementors, and connectors, firms offering adapters or brokering services to combine other firms' products or services. Innovation strategy, then, is very much about choosing a firm's position in the ecosystem in terms of one of these four roles, and employing strategies to move a firm from its current to a desired position. These strategies include (dis)intermediation, (de)modularization and insourcing/outsourcing, changing a firm's internal and external dependencies. Another key result from the application of the ecosystem complexity model is its use to understand the lifecycle of industries. In summary, many new industries are populated by completers, by lack of other firms. Connectors find a role to combine products of different completers, or of composers that start to offer platforms, applying emerging standards, enabling the rise of complementors.

5.1 Introduction

The interest in the notion of business ecosystems has seen a steady increase in the past decades, both in science (Moore, 1993; Adner, 2006; Iansiti & Levien, 2004) and practice (LeHong et al., 2019). Awareness has been growing that organizations do not operate in a vacuum and have many interdependencies with other actors including customers, employees, suppliers, partners, unions, NGOs and governmental organizations. However, the notion of business ecosystem does not stop at this general observation. In fact, there is a growing body of work that provides further insights about the nature, dynamics and importance of ecosystems highlighting the key role of ecosystems in innovation processes (Adner, 2017; Adner & Kapoor, 2010; Brusoni & Prencipe, 2013; Ethiraj & Posen, 2014; Gawer & Cusumano, 2014a; Jacobides et al., 2018; Mäkinen & Dedehayir, 2014; Russell & Smorodinskaya, 2018).

Notwithstanding the growing ecosystem awareness in management science, many organizations do not seem to pay explicit attention to understanding and managing their business ecosystem. In fact, in a recent global market survey (Burton et al., 2017) with 631 organizations, only 28% of participants, when asked about their ecosystem attention, were primarily focused on their 'public' ecosystem, defined as including both known and unknown actors. 29% of participants were focused on their 'private' ecosystem, defined as including only known and directly related actors, like customers and suppliers. The remaining 43% indicated to be primarily focused on their 'internal' ecosystem, equating that with their internal organization. These results indicate that a vast majority of firms do not prioritize their attention to their ecosystem outside their internal organization and immediate surroundings.

The ecosystem blind spot of many firms is a serious roadblock for innovation. The lack of attention for the critical role of the business ecosystem is a cause of innovation failures (Adner, 2012; Cianci, 2014; Iansiti & Levien, 2004). Oftentimes, firms do not realize, or realize too late, that their innovation must go hand in hand with innovations from other firms to succeed. A classic example being the introduction of HDTV by leading firms such as Philips and Sony. This introduction was premature as complementary innovations in TV and film industries took years to complete, after which other TV manufacturing firms had caught up with Philips and Sony (Cianci, 2014). In other cases firms may not be sufficiently aware that their innovation still needs to be adopted by clients or intermediaries in their value chain (Adner, 2012).

In general, business ecosystems are collectives of both non-competing and competing firms, including firms that compete in one context and collaborate in another context. The non-competing collaboration between firms in business ecosystems may serve many potential purposes. These include sales and distribution of services and goods, procurement and supply, (monetary) value exchange, sharing of services and resources, and so on. Here our interest lies with aspects of business ecosystems that are relevant to innovation (Adner, 2006; Iansiti & Levien, 2004; Barnett, 2006).

A key aspect is that effective innovation should not only take the ecosystem as-is into account, but may also require the ecosystem to be reconfigured, specifically its structure of interdependencies between firms and their products and services, in order to support the innovation. One may even argue that reconfiguring or influencing the

architecture of the ecosystem's structure is an essential part of the innovation process itself (Adner & Kapoor, 2010; Ethiraj & Posen, 2014). With the demise of complete, stand-alone products and the rise of platform technologies and their complementary products, an increasing number of scholars have argued that the architecture of the ecosystem dependencies, especially between platform firms and complementors (Brandenburger & Nalebuff, 1997), is a crucial factor for innovation (Tee & Gawer, 2009; Cennamo & Santalo, 2013; Gawer & Cusumano, 2014a; 2014b).

Most extant research in this area has focused on empirical studies and inductive theorizing. A theoretical understanding of innovation ecosystems is still largely lacking. Several scholars have suggested that complex systems theory may serve as a foundation for more fundamental insights in business ecosystems (Peltoniemi & Vuori, 2004; Anderson, 1999; Desai, 2010; Russell & Smorodinskaya, 2018; Luo, 2018). Following these suggestions, the approach in the current work is to apply the perspective of complexity, and see if and how that could add to or alter the insights about innovation strategy in ecosystems and specifically about the position and role of firms in their ecosystem structure.

In this study, we start by examining existing work on the topic at hand. Then, we proceed by adding insights to innovation strategy development by applying the complexity perspective. We conclude with an overview of implications for innovation policy and management.

5.2 Innovation Strategy in Business Ecosystems

5.2.1 Business Ecosystems

The concept of business ecosystem was introduced in 1993 by Moore (1993). As the title "Predators and prey: a new ecology of competition" and the concept of ecosystems indicate, this and many other studies on these topics are inspired by concepts from biology.

Moore defined business ecosystem as (Moore, 1993, p.76):

"An economic community supported by a foundation of interacting organizations and individuals – the organisms of the business world. The economic community produces goods and services of value to customers, who are themselves members of the ecosystem. The member organisms also include suppliers, lead producers, competitors, and other stakeholders."

He then continued with describing how 'member organisms' develop and how they relate to each other over time, which is a key theme in the current work:

"Over time, they coevolve their capabilities and roles, and tend to align themselves with the directions set by one or more central companies. Those companies holding leadership roles may change over time, but the function of ecosystem leader is valued by the community because it enables members to move toward shared visions to align their investments, and to find mutually supportive roles."

It should be noted that business ecosystems do not consist only of commercial firms¹¹, both competing and non-competing (Jacobides, et al., 2018). In fact it is not uncommon that governments are also active ecosystem actors, as well as in some cases universities, public research organizations and NGOs.

Moore' definition of business ecosystem has been refined in other work. For instance, Jacobides et al. (2018) describe ecosystems as a group of interacting organizations that depend on each other's activities. In considering the emergence and dynamics of ecosystems, modularity is emphasized as an enabler of how interdependent organizations coordinate without hierarchical actors. Complementarities or interdependencies are used to explain how ecosystems are different from other supra-firm concepts such as markets and supply chains. If complementarities are not specific but generic (Teece, 2018), and fungible, then there is no strong need for coordination between firms, and interaction can take place through general market transactions. Very specific complementarities require very strong coordination, supported by contracting in supply chains or in-house coordination. In between are complementarities that require coordination in an ecosystem with multiple firms that choose to collaborate, but without strict contracting or formal hierarchical mechanisms. These are complementarities between two products, services or activities in which one requires the other or in which one makes the other more valuable, so-called unique complementarity (Hart & Moore, 1990) or supermodular complementarity (Milgrom & Roberts, 1990), respectively. However, in order for two products to act as complements, coordinated efforts are often required regarding interfacing, standards or anything else that is required without which the two cannot be combined.

Others make a distinction between a more general perspective of business ecosystems (Moore, 1993), and more specific perspectives such as innovation ecosystems (Adner & Kapoor, 2010; Brusoni & Prencipe, 2013), with interactions between firms to co-innovate, or platform ecosystems, with dependencies between platform firms and complementor firms (Kapoor, 2018; Gawer & Cusumano, 2014a; 2014b). Here, we focus specifically on the implications of ecosystems for the innovative activities of firms, rendering the literature on innovation ecosystems especially relevant.

By nature, business ecosystems are not static. Firms are born and die. Most firms that existed 50 years ago are no longer here, with shortening average lifespans (West, 2017), further underpinning the need to think in terms of ecosystems, rather than in terms of an individual firm (Barnett, 2006; Iansiti & Levien, 2004; Adner, 2006). Moreover, connections and dependencies between firms and other ecosystem actors change frequently. That is not to say that all connections have similar purposes. Some connections are focused on sales and supply interactions, others on service provisioning, and still others on co-innovation, or competition. The point here holds that whatever the purpose, these connections are changing frequently, and increasingly so. For example in open markets, more emancipated customers are

¹¹ For (very) large firms and organizations, actors in ecosystems may also refer to company divisions, organizational units or product groups. This implies that different parts of a large company may play different roles in the same or multiple ecosystems. For example, Microsoft's Windows group and Microsoft's Xbox division play very different roles in entirely different business ecosystems.

changing from one supplier to another more easily – the concept of a ‘regular customer’ is in many businesses a thing of the past. And the same goes for business-to-business connections, where increased competition, standardization and other factors have made it easier for companies to for example switch from one supplier to another.

In addition to easier switching of connections, in many cases, it has become easier to initiate and maintain more and more connections. Thanks to advances in communication, mobility, the resulting globalization and especially the advent of the digital age (Dicken, 2015), ecosystems are ever more dynamically and richly interconnected. The result being that the complexity, in the sense of the extent of interdependencies between actors in business ecosystems, continues to grow. It only seems natural, therefore, that the perspective of complexity theory is highly relevant in this context, which we will address in the next section.

5.2.2 Ecosystem Innovation Risks

Most business ecosystems have emerged and developed in an organic, undirected way, based on the evolution of a market or industry, government regulations and other factors. In some cases, however, a business ecosystem may be developed deliberately or strongly influenced by single organization, often a large firm with dominant value chain position.

The latter case is the perspective taken in the work of Adner (2012), who stresses the need for firms to manage two main ecosystem related innovation risks, in addition to a firm’s own internal innovation execution risks:

1. Co-innovation risk: which other firms in the ecosystem need to successfully innovate to make the firm’s own innovation successful? For example, which parts that are supplied by external parties need to be improved in order to make the firm’s new product work? Which complementary products, services or content are needed to make the new product useful to customers?
2. Adoption chain risk: which other actors in the ecosystem that stand in the value chain between the focal firm and end consumers need to adopt the firm’s new innovation? For example, which intermediaries, wholesale or retail actors need to be convinced, on the basis of which value incentives for them?

The success or failure of many innovations can be better understood by applying the above risks for analysis, demonstrating the necessity for firms to look beyond their internal execution and develop innovation strategies that take the business ecosystem perspective fully into account. For example as mentioned before, in the film industry (Adner, 2012), the introduction of digital film as new innovative medium with many quality and cost related benefits, required the co-innovation of different actors in the film ecosystem, including studios and manufacturers of editing, camera and microphone tools and equipment. However, it was especially in the process of adoption that the new digital film took a long time to become successful. Although directors and production companies quickly could benefit from this innovation, this was less true for many cinemas, which had to invest in new and expensive equipment, with little incentive to do so. Ultimately this adoption chain risk was mitigated by the creation of financial intermediaries, funded by large players in the film industry, to facilitate cinemas in making the required investments in an economically viable way.

5.2.3 Innovation Strategy and Ecosystem Position

A first step to develop a more effective 'ecosystem aware' innovation strategy is to understand and choose the position that the own firm and other actors should have in the ecosystem, in order to optimize support for the innovation (Adner, 2012; Adner, 2017; Iansiti & Levien, 2004). To do so, one needs to understand and map the current business ecosystem, and then identify which new or existing roles the own firm and others would need to play in the ecosystem, and which relationships are required, where and when.

According to Iansiti and Levien (2004) ecosystem health can be measured in terms of productivity, robustness and niche creation. Productivity represents the ability of firms to transform inputs like for example raw materials or technological knowledge to outputs such as finished goods or new innovations. Robustness is the capability of the business ecosystem to survive changes and disruptions, like a drop in the availability of natural resources or unexpected technological change. Niche creation refers to the ecosystem's capacity to support sufficient variety, for example in applying emerging technologies in multiple existing and new products, supporting value creation and also robustness.

Firms can take on different roles in ecosystems. Iansiti and Levien (2004) make a distinction between three different ecosystem roles: keystones, dominators and niche players. Keystone firms in ecosystems are crucial hubs in regulating the ecosystem health. Not just of the individual firms that play the keystone role, but of all firms that have interdependencies with the keystone firm. Keystone firms regulate these ecosystem health measures (productivity, robustness, niche creation) by providing other firms with a set of commonly useful assets. Although typically small in number, keystone firms provide a foundation or platform of assets and capabilities that supports or improves the productivity of a diversity of many other firms or 'niche players' that focus on specific products and services. Moreover, a stable and predictable keystone platform helps firms in their capability in responding to changes, improving the robustness of the ecosystem.

As a result of their pivotal role in ecosystems, keystone firms have a pervasive impact. Nevertheless, keystone firms still depend on the health of many and often smaller niche players. The keystone firm and associated niche players are complementary: both are needed to create value and provide end-customers with useful products. If keystone firms fail to recognize this mutual interdependency, and start to overuse the power that comes with their position, then their role changes into what is known as a 'dominator' (Iansiti & Levien, 2004). A dominator may play its role in two ways: either to squeeze out the value of the rest of the ecosystem by for example lowering profit margins of niche players by charging higher prices for the use of keystone facilities, or by acquiring niche player firms in the ecosystem, thus directly controlling and harvesting the value creation in ever greater parts of the ecosystem. Both dominator strategies tend to be short-lived by severely degrading ecosystem health: either by loss of productivity and robustness of niche players, or by loss of diversity and robustness, by limiting the number of firms in the ecosystem.

A well-known example of a keystone firm is Microsoft, or rather its Windows product group (Iansiti & Levien, 2004). The Windows operating system for PCs and other computing devices, including accompanying configuration and development tools,

has a large majority market share and is used by myriad niche players to develop business and personal applications. The value creation in this ecosystem is of enormous scale, supporting the productivity of, amongst others, many software developers. The ecosystem also benefits from a reasonably predictable platform with backward compatibilities, providing a stable environment for niche players to do innovation. In fact the success of this ecosystem has grown to such proportions that there have been growing concerns about the uncontrolled power and influence and risk of dominator behaviour by Microsoft and other similar companies like Google and Apple (Barnett, 2006).

5.2.4 Innovation Strategy & Ecosystem Reconfiguration

As part of a more effective innovation strategy, it might be necessary for a firm not to accept its ecosystem as it is, but to reconfigure the ecosystem or build a new ecosystem entirely. Adner (2012) illustrates this by describing examples of innovation failures and successes. In his view, innovation success depends on firms consciously taking a leading or following position in their ecosystem. Depending on this position, firms either actively involve other firms or more passively wait for other firms, ensuring that time and place of new innovations are synchronized with the readiness of other ecosystem firms that are needed for co-innovation or adoption. Going further, leading firms may pursue a strategy to reconfigure their ecosystem in order to maximize innovation success. Reconfiguration occurs when the leading firm introduces new firms such as for example intermediaries or platform providers to support other firms in co-innovation or adoption, or when leading firms acquire other firms to have more control on required co-innovation and adoption, in both cases changing the structure of the ecosystem. An example of which can be found in the abovementioned case of the film industry, with the introduction of financial intermediaries to facilitate adoption of digital film by cinemas.

The perspective of Iansiti and Levien (2004) on innovation strategy is focused on how a firm can choose or change its role, depending on the characteristics of their ecosystem, and how that in turn may affect their ecosystem. They argue that in dynamic and innovative ecosystems, a firm can be successful by either choosing a niche player or keystone position. As a niche player, firms focus on their own unique area of expertise to differentiate themselves from competitors. This focus and their leveraging of more or less stable assets provided by keystone actors, allows them to survive in the turbulence of their ecosystems. Keystone firms seek a position in the ecosystem network that allows them to make their assets or platform available and useful for (many) other niche player firms. If the keystone firm is able to do this in a cost efficient and scalable manner, then it will be able to capitalize on its influential position by sharing in the value creation that it facilitates in many other firms. Using this approach, the keystone will also foster diversity by lowering the hurdle for new innovations and new firms, and thus increase the robustness of its ecosystem to respond to changes. If the latter is of less importance, in a more stable and less innovative ecosystem, a firm may seek to become a dominator, which allows a firm to maximize its share of harvesting from the value creation in the ecosystem, provided that this strategy does not undermine the ecosystem health, in which case the dominator strategy will be short lived. Another position for a firm in a more stable

ecosystem is to focus on its own niche, which however will be more of a commodity, and therefore likely less profitable, more vulnerable to competition and to ultimately unavoidable changes in the ecosystem equilibrium.

Adner and Kapoor (2010) distinguish between challenges for a firm arising from innovations occurring upstream or downstream in a value chain. They refer to this by the notion of technological interdependence. If firms are confronted by an upstream innovation, such as a component innovation, they may be forced to change their own technology to accommodate this change and seize the opportunities provided by a new component technology. Hence, in such cases performance improvement through learning is more likely as a firm is forced to change its status quo. Thus, the complex interdependencies provide scope for learning, and successful co-innovations in such a complex context also render it harder for competitors to do imitation (Rivkin, 2000). However, this is tempered by increases in modularity (Baldwin & Clark, 2000), as complexities inside modularized components become less relevant to focal firms which can more easily swap one component supplier for another. Modularization reduces complexity or, in terms of Adner and Kapoor (2010), reduces the challenge, learning opportunity and performance improvement potential. By contrast, if firms meet downstream challenges when confronted with complements that are not yet available or take longer to develop, then this will delay and may also erode the performance improvement and value creation potential in the focal firm. Not only does a delay offer competitors a chance to catch up, it also reduces the focal firm's opportunity for learning, because production and possible further practical improvements from i.e. usage feedback are not yet happening.

Taken together, upstream challenges enhance while downstream challenges erode the competitive advantage of early movers or technology leaders. Adner and Kapoor (2010) find evidence for their claims from an empirical study on decades of innovation in the global semiconductor lithography equipment industry. However, it seems that these arguments are weakened by a number of notions, such as the consideration that delays in component innovation and supply have the same eroding effect as complement delays. Likewise, complements that require adoption of the focal solution in order to have them work together, also creates learning opportunities - learning is not a unidirectional affair.

In addition, Adner and Kapoor (2010) argue that the effectiveness of the strategy of vertical integration, to manage or influence technological interdependencies and therefore complexity, changes during the lifecycle of the technology at hand. Early in the lifecycle, technology challenges tend to be largest, while decreasing over time with more innovations being successful in addressing these challenges. Therefore, vertical integration of multiple components into the focal firm during early phases only serves to move this complexity from the external to the internal, overshadowing the benefits of vertical integration in reduced contractual or supply chain issues. However, with technological complexity decreasing over time, the benefits of vertical integration start to count more.

In a more recent contribution Kapoor and Agarwal (2017) introduce a framework that focuses on the performance of complementors rather than on the keystone firm. This framework is mainly based on the notion of complexity in terms of interdependencies between a complementor and the keystone firm whose platform it is working with.

The framework was validated in an empirical study on the performance of app builders, which complement Apple's iOS or Google's Android platform, or both. One of the findings was that complementors with greater internal and/or external complexity are hard to imitate by other complementors, protecting the unique competitive position of the focal complementor and sustaining its performance level. In addition, they also recognized that because of interdependencies, changes in one firm may affect the performance of other interdependent firms. Or in their more specific case, they find that changes in the keystone firm or its platform make it harder for complementor to sustain their current performance levels. Indeed, if complementors work with multiple platform providers, then too frequent platform changes may become overwhelming for an app builder. They conclude that high complexity is good to protect against imitation, but is also making it harder to deal with platform changes and the resulting changes in the fitness landscape.

Finally, the ecosystem perspective also relates to earlier work by Christensen (1997), who stresses the fact that firms in their innovation depend on other firms in their value chain, especially customers. One of his key points is that newly innovated products often require new customers, at least initially, as existing customers have a tendency to prefer existing products, meeting their current requirements, and to resist new, unfamiliar products. To innovate successfully, firms also need to work on their connections in their value chain, building up a new customer base. Nevertheless, most attention in Christensen's work still seems to go to improving innovation management principles within the firm, with the risk of insufficient managerial attention for the external business ecosystem.

5.3 Complexity Perspective on Business Ecosystems

Ecosystem scholars agree that developing an effective innovation strategy requires an understanding of a firm's position in its ecosystem and its connections and interdependencies with other ecosystem actors (Adner, 2006; 2012; 2017; Kapoor, 2018; Kapoor & Agarwal, 2017; Iansiti & Levien, 2004). Taking this as a starting point, the goals of any innovation ecosystem strategy are to (a) determine the most optimal position for the firm, given its innovation ambitions, and then (b) find ways to move the firm from its current to its desired position in the ecosystem.

To shed more theoretical light on the insight gained from ecosystem studies, we turn to complexity theory. This perspective can be deemed relevant as the sheer notion of an ecosystem refers to a system of interdependent parts, where the parts refer here to firms and the interdependencies to complementarities between their products or services. In particular, we will make use of the NKCS model developed by Kauffman and Johnsen (1991), which models interdependencies between co-evolving parts in an abstract and formal manner, allowing it to be transferred to the realm of business ecosystems (for a review, see Vidgen & Bull, 2011).

5.3.1 Internal and External Complexity

To understand a firm's ecosystem, two complexity perspectives are relevant (Kauffman & Johnsen, 1991). The first perspective of *internal* complexity, or the extent of dependencies between units or components within a firm, is relevant to understand

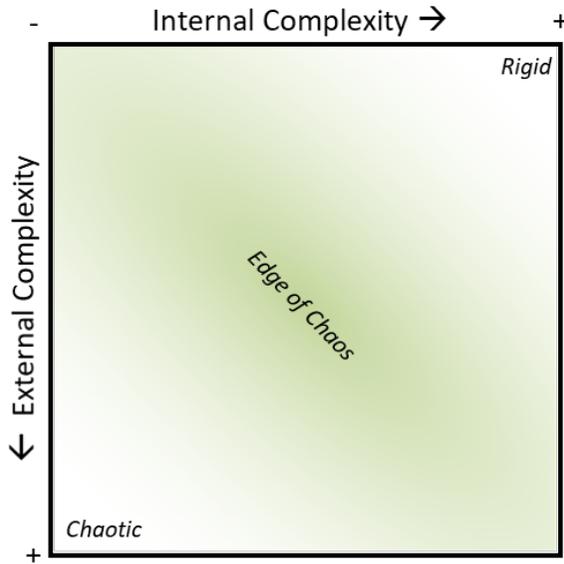


Figure 5-1. Ecosystem complexity model, characterizing the ecosystem regime by the internal and external complexity levels. The intensity of the color is a representation of the relative level of performance.

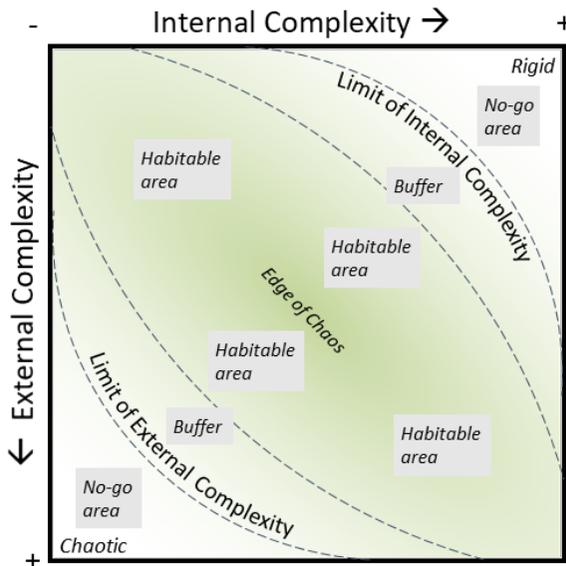


Figure 5-2. Different areas in the ecosystem complexity model, indicating where firms are able to innovate (and where not), depending on the (mis)balancing of internal and external complexity.

a firm's agility to change and to innovate. The second perspective of *external complexity*, or the extent of interdependencies between firms and other actors in an ecosystem, is relevant to understand the process of co-innovation between firms. If one firm makes a move such as innovating a product, then other dependent and therefore impacted firms have to respond by doing some innovation on their own, and so on (Schumpeter, 1939; Silverberg & Verspagen, 2005; Arthur, 2007). This co-innovation of multiple interdependent firms may be either deliberately and planned, or reactive and unplanned, emerging from the exchange of products and services. But either way, the dynamics of co-innovation may differ greatly between ecosystems (Kauffman & Johnsen, 1991).

On the one hand, with more dependencies between more actors, and hence with more external complexity, the process of co-innovation typically becomes more intense and takes longer. The process may even become chaotic in the sense that a single innovation ignites an endless chain of innovation responses, which may trigger yet other innovation responses by other actors, even back to the originating actors (Luo, 2018). Or, formally, a single innovation may upset a state of Nash equilibrium (Nash, 1951) in such a way that the subsequent series of responses never settles into a new Nash equilibrium, for which holds that no single firm can improve its fitness by mutating. Once in such a chaotic regime, it is very hard for firms to improve their performance by innovation, as their performance depends on so many other firms that are innovating and causing changes so often that it becomes very hard or even infeasible to book any sustainable progress. For example in the current open source software market, with many incumbent and start-up firms and individual actors working together in frequently changing consortia, new innovations occur so often that firms in this ecosystem, including investors and customers, continuously have to reconsider and refactor in order to keep up (Shah, 2006), in some cases to such an extent that products become obsolete before even reaching production stage (West & Gallagher, 2006).

On the other hand, a sparsity of connections or dependencies between firms inhibits the process of co-innovation in search for complementarities (Adner & Kapoor, 2010; Luo, 2018). One firm's innovation is not leading to responsive innovation in other firms, if there are no dependencies between firms. In other words, in contrast with the chaotic regime, a low level of external complexity may ultimately lead to a static or rigid innovation regime as firms are hardly triggered by others firms to innovate. An example of this would be a firm that develops and sells its own commodity products to a few clients, having a very limited number of suppliers who deliver raw materials that are also commodities. In such a situation, a firm is less likely to find justification to invest and spend resources on innovation, thus limiting their progress in performance. For example, traditional mechanical watch makers operate in a rather limited and static ecosystem, with a limited number of suppliers and (wealthy) customers, and watch makers doing most of the construction work themselves, using their own specialized knowledge and experience, with a relatively low rate of innovation. Interestingly, even the advent of digital and smart watches has not significantly changed their way of working, although it did greatly reduce their market share and pushed traditional watch makers to a more exclusive market niche (Norman & Verganti, 2014).

However, the dynamics of the co-innovation process in ecosystems are not only dictated by external complexity. The ability of firms to do innovation themselves, is also highly relevant. From a complexity perspective, internal innovation is also to be understood as guided by dependencies between the different components that make up the subject of innovation, like a product or service. This internal complexity (Kauffman & Johnsen, 1991), like external complexity, is an important determinant in the innovation regime of the ecosystem. If internal complexity is high, then firms find it harder to innovate, thus moving the regime towards a more rigid state. And *vice versa*, with lower internal complexity, firms can innovate more easily, moving the regime towards a less rigid or more chaotic regime. The latter may occur, amongst other causes, when firms start to modularize their solutions (Baldwin & Clark, 2000), in effect lowering the dependencies between constituent parts (Fleming & Sorenson, 2001a; Frenken & Mendritzki, 2012).

Taken together, the interplay of internal and external complexity are major factors in driving the ecosystem towards either chaotic or rigid regimes. In both extremes, performance suffers, either from too much change that renders improvements quickly useless, or from a lack of change that inhibits progress. It follows that the optimal regime, identified by Kauffman and Johnsen (1991) as 'the edge of chaos', is where the ecosystem offers enough dynamics for frequent improvement, yet avoiding too much chaos and endless and destabilizing avalanches of innovation responses. In other words, firms that balance internal and external complexity levels find themselves in the 'edge of chaos' regime that provides both enough stability to reach higher levels of performance, and enough adaptability to maintain or restore performance levels in case of changes. Using simulations, the edge of chaos regime is found to not only generate the highest performance levels, but also the best robustness characteristics in terms of performance restoration after ecosystem shocks (Kauffman & Johnsen, 1991; Vidgen & Bull, 2011). See also chapter 4 in this thesis.

5.3.2 Ecosystem complexity model

The internal and external complexity dimensions are used here to create a model to characterize different ecosystem regimes, as shown in figure 5-1. Internal complexity, from low to high, is represented in the horizontal dimension. Likewise, external complexity is represented in the vertical dimension. Using these two dimensions, three ecosystem regimes are indicated: rigid, chaotic and the edge of chaos. The rigid regime is characterized by high internal complexity, with firms finding it hard to change their fine-tuned configuration of components without losing performance, and by low external complexity, with firms having little need to change as they hardly depend on other firms. The chaotic regime is characterized by high external complexity, with firms having to respond to many changes in many other firms on which they depend, and by low internal complexity, with firms also having the internal flexibility to respond to change. The edge of chaos is characterized by a balance of internal and external complexity, avoiding the rigid regime with lock-ins on suboptimal performance levels that are hard to change, as well as avoiding the chaotic regime with very frequent changes that render improvements useless even before completion. The color intensity in the figure 5-1 reflects the qualitative performance differences between the different regimes. Performance levels are highest at the edge of chaos, and degrade

with more unbalanced internal and external complexity, leading to either more rigidity or more chaos (Kauffman & Johnsen, 1991)(see also chapter 4 in this thesis).

More subtly, performance is also higher at intermediate complexity levels. This results from the fact that products and services with low complexity typically consist of commoditized parts with performance being simply the sum of its parts as complementarities are absent. At high complexity, the interplay of parts requires much more fine-tuning, which – if successful – results in a performance of the whole that is more than the sum of its parts, reflecting the many complementarities. However, given the many interdependencies an optimal configuration is hard to find, with a minor change likely resulting in a sudden performance drop, implying that on average, performance will be relatively low when complexity is high. The best performance, on average, is found at intermediate complexity levels, avoiding both the commodity of low complexity and the sensitive fine tuning of high complexity (Kauffman & Johnsen, 1991). See also chapter 4 in this thesis.

Taken together, performance is highest at the edge of chaos and at intermediate complexity level, resulting in the performance levels that are depicted in figure 5-1.

Within the ecosystem complexity model, different areas can be distinguished in terms of the extent to which firms would be able to survive in the different regimes of the innovation ecosystem. See figure 5-2. The most extreme regimes, either very chaotic or very rigid, are practically “no-go areas” for firms to innovate. Both extremes are characterized by a strong misbalance between internal and external complexity. These no-go areas exist where a firm has reached either its limit of external complexity, relative to internal complexity, or its limit of internal complexity, relative to external complexity. Habitable zones are found where internal and external complexity are more balanced, on or around the edge of chaos. In between habitable and no-go areas are buffer zones, in which firms still may be able to innovate, but are at risk of moving too far away from their habitable zone.

5.4 Ecosystem positions

The ecosystem complexity model is primarily used here to characterize the position of a firm in its ecosystem, either for the entire firm if the firm’s portfolio is not too heterogeneous, or for an organizational unit that is responsible for one or more particular artefacts, like specific products or services under consideration. To determine a firm’s or unit’s position, a qualitative estimation of its internal and external complexity levels is required. For internal complexity, this may be done by answering questions like: if one part of our product is changed, which percentage of other parts also have to change? Or: if one team changes its way of working, which percentage of other teams have to change as well? For external complexity, this may be done by answering questions like: which percentage of the parts in our product depend on external suppliers? If we were to innovate our product, how many of our clients would need to adapt their use of our product?

Other ways to gauge a firm’s position in the ecosystem complexity model is to consider the regime in which the firm finds itself. To what extent are long chains of related innovations perceived, in which the firm takes part? To what extent is the

firm triggered by other actors to do innovation? How easy or difficult is innovation experienced within the firm?

Furthermore, the position of firms in the ecosystem complexity model can be organized in 4 categories, inspired by Berends et al. (2018), that result from combining low and high levels of internal and external complexity, as shown in figure 5-3:

- *Composers*: firms that enable and facilitate the combination of products or services from other firms with their own products or services, often provided as a platform or infrastructure for others to use or to build on. Often, these firms (or parts thereof) have to be able to deal with relatively high complexity, as their platforms or infrastructure are made up of many components, with both internal and external interdependencies. This complexity may be reduced by modularization (Baldwin & Clark, 2000; Fleming & Sorenson, 2001a; Frenken & Mendritzki, 2012), to some extent. In a sense, the added value provided by composers is to 'hide' their complexity from other firms in the ecosystem, who can then focus on complementing the offering of the composer with their own less complex services or products, which in turn renders the offering of the composer more relevant and increases its market share. An example of a composer firm is Google with their Android operating system for smart phones, including development tools and their 'playstore' as a market place, offering a platform to numerous app developers.
- *Completers*: firms that seek to innovate their own, often unique products or services, complete in their own right, with limited or no dependence on other firms. Such firms have high internal complexity, as all required and interdependent parts (often commodity parts to limit supplier dependence) are carefully synthesized by the firm itself. Typically such firms possess specialized knowledge to design and manufacture their product, which is hard to copy by competitors (and if not, then completers will have a hard time to survive). Completer firms strive to maximize autonomy and control by minimizing dependency on other (exchangeable) firms, resulting in lower external complexity. An example of a completer firm is IBM in the 1960s, producing and controlling all hardware, software components and services that comprised their mainframe computer offering. Later, IBM changed their position with the advent of mini and microcomputers, amongst others, with open standards that allowed the entry of other firms to provide components (Baldwin & Clark, 2000), at which point IBM became more like a composer firm.
- *Connectors*: firms that add value by offering products or services that facilitate the connection between products or services from other firms. Typically these firms limit their internal complexity by adopting open or industry standards, rather than dictating their own standards to others, and by keeping their own product or service (as a connector between others) relatively small. However, the result is that, if needed, such firms must be compliant with multiple standards, leading to higher levels of external complexity. An example of a connector firm is Wink (Berends et al., 2018), which enables the working together of home automation products and platforms from different vendors, including vendors like Apple, Amazon and Google, many of which use different open or proprietary standards.

- *Complementors*: firms that develop their own products or services that are often built on top or are combined with products and services from other firms, typically infrastructure or platforms provided by composers (Kapoor, 2018; Gawer & Cusumano, 2014a; 2014b). Complementors work with a limited number of other (composer) firms or only a single composer, limiting their external complexity. By leveraging the often standardized and modularized offerings from other firms as a foundation for their own products and services, complementors also have a relatively limited internal complexity. However, the latter should not be too low, to make their products or services differentiate sufficiently from competitors, avoiding commoditization. Examples of complementor firms are abundant, as they make up the bulk of firms in most ecosystems. Think for example of Adobe selling document viewing and editing software, built on top the Microsoft platform, or for example EnergieDirect, selling energy to consumers in the Netherlands, depending on infrastructure and energy production from other energy firms.

The ecosystem complexity model can be used to indicate the current and relative position of a single or multiple firms in their ecosystem. As an example, figure 5-4 indicates the position of several firms in the ecosystem of mobile phones & tablets. A telecommunications company such as Vodafone is a clear example of a composer, providing the infrastructure and services that enable all other firms in this ecosystem to develop and sell their products and services. Other examples of composers are Apple and Google, both providing a mobile device operating system (iOS and Android) and other composer services, such as a market place for apps (Appstore and Play). As Apple is also a provider of hardware like iPhones and iPads, this firm is closer to the completer quadrant than Google. Examples of the very many complementor firms in this ecosystem include smart phone manufacturers and mobile app providers, such as Samsung (Galaxy phones) and Willowtree (mobile apps). Samsung utilizes the Google Android platform and is therefore not considered to be a composer. But Samsung is relatively close to the completer quadrant as it is also producing its own apps and also offers accessories like smart watches, not to mention many other Samsung products. Willowtree, on the other hand, is closer to the connector quadrant, as it develops apps both for iOS and Android. An example in this ecosystem of a true connector is Citrix, which provides products that allow apps to run on multiple platforms and operating systems, including cloud platforms, Windows, Linux, iOS and Android. An example in the opposing quadrant of completers is BlackBerry, as far as it is still in business. This firm was one of the first in this ecosystem, and offered a complete service that included an early version of the smart phone, an operating system and several apps. Only when composers such as Google and Apple entered, BlackBerry gradually started to move to become a composer too, by amongst others opening its own app market place, but not before it already had lost most of its market share (Querbes & Frenken, 2017).

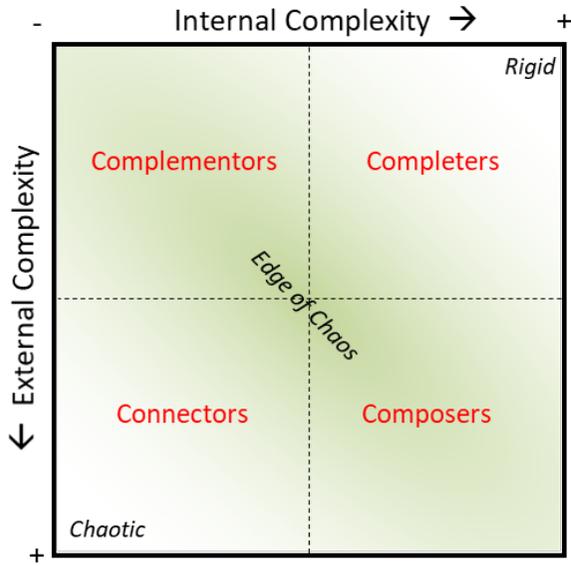


Figure 5-3. Ecosystem complexity model with four firm position categories.

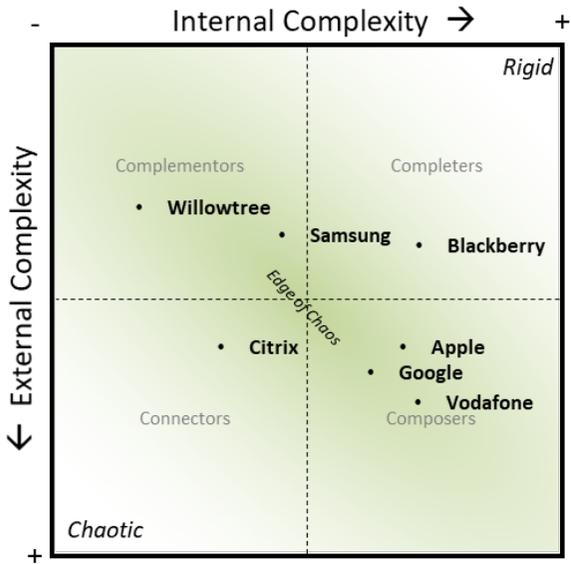


Figure 5-4. Example: position of a number of firms in the mobile phones & tablets ecosystem.

We should note that the ecosystem complexity model and its position categories are relative to a specific economic perspective or context. Firms may be a completer in one perspective, while being a composer or complementor in another. For example, an airport may be a composer in the perspective of air transportation, as provider of infrastructure and literally a platform that is complemented by airline companies, air transport service providers and so on. Airline companies in turn may be considered as composer in the sense that they orchestrate the working together of pilots, catering, fuel suppliers, airplane maintenance providers and airplane manufacturers. The latter, in turn, may be considered as composer in the airplane manufacturing perspective, in which case they provide the architecture and orchestration as a platform to bring together the components of countless part manufacturers that act as complementors in this perspective.

5.4.1 Ecosystem positions and relations

From an ecosystem perspective, firms should not be considered in isolation (Adner, 2012). Amongst other characteristics, a firm's position category does not stand alone, but must be in symbiosis with the position category of other firms (Brusoni & Prencipe, 2013; Gawer & Cusumano, 2014a; Kapoor, 2018). For example, composers can only exist if there are also complementors. Connectors typically only flourish in the presence of multiple composers, as will be discussed further below.

Considering the position of the edge of chaos in which innovation circumstances and performance levels are most optimal (Kauffman & Johnsen, 1991), it becomes clear that composers and complementors are more likely to be close to the edge of chaos, as they both tend to balance their internal and external complexity, avoiding combinations of low internal and high external complexity, or *vice versa*. Typically, complementors do this by relying on the offerings of one or a limited number of composers (Kapoor & Agarwal, 2017), which they use to reduce the internal complexity of their own products, while still keeping their external dependencies in check. Composers do this by differentiating themselves with platforms or infrastructure that are reliable, rich in functionality and scalable enough (Boudreau, 2010; Gawer & Henderson, 2007; Gawer & Cusumano, 2014a; 2014b; McIntyre & Subramaniam, 2009) to be used by many complementors, thus combining their sophisticated offering with a dependency on the products of complementors, without which their platform or infrastructure would be barren and therefore useless.

This symbiosis of composers and complementors has become ever more pervasive in many ecosystems in markets like IT, telecommunications, transportation, tourism and others, leading to what is also known as the 'platform economy' (Kenney & Zysman, 2016). In all cases, globalization and advances in digital technology and communication, including the Internet, have helped firms (Dicken, 2015) and in particular composers in their ability to scale their offerings virtually unlimited and manage their symbiotic networks with many complementors cost efficiently (Cusumano & Gawer, 2002; Cennamo & Santalo, 2013). In addition, the collaboration or compatibility between composers and complementors is very much facilitated by the availability of open or proprietary (i.e. composer-owned) standards (Gawer & Cusumano, 2014a; 2014b). Other factors that explain the rise of composers and complementors include political and economic tendencies to open up markets to foster more competition

and innovation (Markard & Truffer, 2006), basically breaking up completer firms and making room for composers and complementors. Examples are the break-up of (state owned) rail transportation (Kleinová, 2016) and energy companies (Joskow, 2008) into infrastructure, service or commercial firms, also opening up the market for new complementor entrants.

Cennamo and Santalo (2013) point out that composers that provide platforms, may often benefit from network effects, meaning that consumers will be attracted to platforms with the most complementor products and services, or the most other consumers, which in turn attracts more and new complementor firms, further enriching the number of platform complements and attracting even more consumers. For example, a main driver for video gamers to select the platform of their choice, such as Playstation or Xbox, is the number of games to their liking that is available (Rochet & Tirole, 2003). Likewise, consumers will choose social media platforms that have the greatest number of users, preferably from their existing social context. These network effects predict a winner-takes-all outcome of market competition, with one or only a few composers becoming dominant, which is indeed what can be observed in several industries, including video gaming, cloud computing or social media. Interestingly, however, a complete dominance of a single composer, despite its efforts to attract as many consumers as possible and by stimulating and facilitating the innovation of complements, may be affected by complementors that choose and manage to work with multiple platforms, not wanting to put all of their eggs in the same basket. Also, network effects may have limited (geographical) scope (Lee et al., 2006) and different consumer groups may have different preferences (Armstrong & Wright, 2007), leaving space in the market for multiple composers.

Importantly, composers must be careful in their strategy to not discourage complementors to innovate (Yoffie & Kwak, 2006). Existing work (Cennamo & Santalo, 2013; Gawer & Cusumano, 2008; Yoffie & Kwak, 2006; Gawer & Cusumano, 2014a; 2014b) on platforms, which we understand as typical composers, indicate that because of network effects, the stimulation and facilitation of innovation by complementors is at least as important as the attraction of consumers through sales and marketing. In other words, the market for composers is two-sided: on the one hand consumers, on the other hand complementors (Rochet & Tirole, 2003; Armstrong & Wright, 2007). This provides further underpinning for the critical need for ecosystem-awareness in firms. Recommendations for composers to manage their relations with their ecosystem complementors, include technical support, pricing and marketing as incentives for complementors to select the focal composer, yet balanced with sufficient diversity and competition (Boudreau & Lakhani, 2009) amongst complementors to fuel innovation to create and improve the quality and range of complements. Composers may also try to claim exclusivity for complements on their platform to attract more consumers, but must balance this with sufficient incentives for complementors to accept exclusivity, such as the ease of development and support of complements, advantageous pricing, revenue sharing and sales & marketing support to get access to large and interesting consumer groups. Composers that fail to address their dependencies with complementors effectively, may struggle to attract complementors to their platform. Or they may attract complementors that are only willing to invest in low quality complements if the composer demands

exclusivity, does not offer attractive enough incentives or when competition amongst complementors on the same platform is experienced as too fierce.

In contrast to composers and complementors, connectors and completers seem to be more vulnerable to be too close to respectively a chaotic and rigid regime, unless they manage to position themselves closer to the edge of chaos regime. Connectors may find themselves overwhelmed by external dependencies if many other firms that they try to connect with, start innovating and hence changing their products or standards. Connector firms need to find a way to adapt quickly to multiple standards to yield sufficient value and performance. They will thrive when they manage to combine products and services from other firms that would otherwise not be compatible, for example by creating adapters or multi-purpose interfaces. However, when connector products or services become too (internally) complex, it will become harder and harder to keep up with changes in external products or standards.

Completers, on the other hand, may find themselves locked-in on their self-reached local optimum – reaching a higher level becomes ever more difficult with increasing internal complexity. Sooner or later their performance is bound to degrade, as a result of new technology or new competition, at which point the completer may not be able to respond, by lack of resources or new knowledge, to find a new optimum at sufficient level. Nevertheless, completers may still be in good shape if they manage to achieve a sufficient performance level, and if changes in their ecosystems are not too frequent or disruptive.

The position categories as described above, resulting from the complexity perspective, enrich existing insights about firm positions in ecosystems, most notably the distinction being made between between keystone, dominator and niche player roles (Iansiti & Levien, 2004). However, since this distinction was made from a different perspective which was uninformed by complexity theory, there is no reason to assume a one-to-one correspondence between these roles and the four categories here. At first sight, the keystone role seems to be closest to the composer category. However, if the keystone firm decreases its internal complexity by for example outsourcing parts of its business (Celo et al., 2018), the keystone may also be closer to the connector role. For instance the companies Microsoft and Li & Fung are both described by Iansiti and Levien (2004) as keystone firms. Microsoft plays its role by providing its niche players with assets like software development tools and an operating system. Li & Fung, a large intermediary firm in the apparel industry that brings together demand from retailers with supply from manufacturers, does not provide assets but plays a role by managing many relations and facilitating their collaboration. Microsoft is clearly a composer, whereas Li & Fung is clearly a connector, yet both play keystone roles. The niche player role appears to be closest to the complementor role, but if the niche player increases its internal complexity, possibly by doing more internally and less depending on composer firms, then a niche player may also qualify as a completer. The completer category may also contain firms that play a dominator role, assuming that such firms play the role of dominator by acquiring other firms, thus having more internal and less external complexity. But a dominator that plays its role as a keystone that squeezes out its dependent niche players may still qualify as a composer, albeit less friendly and probably short lived. In other words and in short, the current categorization provides a richer and more foundational explanation than

'dominators' or other ecosystem position characterizations, by taking the complexity perspective into account.

5.5 Moving in the Ecosystem

Determining a firm's current position in its ecosystem is just a starting point in developing a firm's ecosystem-driven innovation strategy. In order to support innovation goals, maximize performance or minimize innovation failure risks, a firm's current position may not be the best place to be, and other positions may be preferable. Using the ecosystem complexity model, firms should strive towards a position that is as close as possible to the edge of chaos, using one of the four position categories as additional direction guide.

After having established a newly desired position, an innovation strategy would clarify how a firm could move from its current to its desired position. From a complexity perspective this means that a firm should initiate strategies and actions that change the levels of internal and external complexity. Baum (1999) examines a variety of such strategies, including breaking up a large company in smaller firms, the redesign of a firm's internal structure to lower (or increase) dependencies between units, or the synchronization of performance measures and incentive schemes to promote collaboration and likely to increase internal complexity. However, these strategies are not limited to innovation also relate to other firm activities, such as production and sales. Here we focus on three strategies that seem to relate most directly to innovation performance: (dis)intermediation, in- or outsourcing, and (de)modularization.

5.5.1 Complexity Strategies

One way to change primarily external complexity is to introduce or remove intermediaries between a firm and other ecosystem actors (Adner, 2012). See figure 5-5. Intermediaries typically reduce the number of clients, suppliers or other actors that a firm has to interact with directly, thus reducing the external complexity of the focal firm. Disintermediation, by for example embarking on direct sales through the Internet, instead of indirect sales through local retail firms, has the opposite effect. To a lesser extent, on average, it is likely that intermediation also reduces internal complexity somewhat, because a firm needs to do less by itself, and *vice versa*.

Other strategies that impact complexity are, amongst others, outsourcing and insourcing (or mergers & acquisitions) (Frenken, 2006a; Celo et al., 2018). Outsourcing can be achieved by moving the innovation, production or management of one or more components to another firm. This may also occur when a firm sells off departments or business units related to one or more components. Either way, from a complexity perspective this means that the focal firm reduces¹² its internal complexity, and increases its external complexity. The reverse happens with insourcing and mergers & acquisitions. See figure 5-6.

¹² There may be exceptions when a firm happens to outsource components that are already modularized, for example, while other components that remain in the firm are not modularized and highly interdependent. In such a case, internal complexity may not decrease, but increase, relative to the number of components within the firm.

Another strategy to change internal and external complexity is the architectural innovation of a product or service (Frenken, 2006a; Frenken & Mendritzki, 2012; Henderson & Clark, 1990). For example, increasing the modular character of a product architecture, is essentially the same as reducing the number of interdependencies between components. Modularization (Baldwin & Clark, 2000; Fleming & Sorenson, 2001a) is closely related to the adoption of standards that facilitate the ease of combining or interfacing between different components (Frenken & Mendritzki, 2012). Instead of direct interfacing between components, modularization entails the introduction of a common interface as a new component, to which other components connect using a standard interface. Standards, typically agreed at industry level or imposed by governments, allow companies to treat components more like a black box, or module, reducing interdependencies. Consequently, components become more interchangeable, lessening external dependencies too if these parts are externally supplied. The reverse effect occurs when components become more synthesized or de-modularised, increasing their interdependencies or complexity, typically requiring more tailored, non-standard interfacing. See figure 5-7.

5.5.2 Mover Typology and Patterns

Following again the ecosystem complexity model, the above strategies can be classified in several types of 'movers', ordered in pairs of opposite directions, see table 5-1 and figure 5-8. The term 'mover' is used to emphasize that applying one or more of these strategies will move or change the position of a firm in its ecosystem. The effect or end-position of each move obviously depending on its starting position and the extent of its impact on internal and external complexity. Of course, firms may also apply different strategies simultaneously, combining their effects to move the firm in the ecosystem. To explain this, a few examples are introduced:

- A A completer may move its position more towards the edge of chaos by outsourcing one or more of its activities. If the firm is in a very rigid (right upper corner) starting point, then only limited outsourcing may still leave the firm in the completer quadrant. If the completer firm's starting position is already closer to the edge of chaos, then outsourcing many components to a single other firm will likely move the firm towards a complementor position. In contrast, outsourcing just a few parts to different other (complementor) firms will move the completer firm likely to a composer position.
- B If the focal firm would be a connector as starting point, then outsourcing would move the firm only more to a chaotic regime, in the lower left corner of the model. Unless the firm is already in that extreme position, in which case there is probably not much left to outsource.
- C A complementor firm that works with multiple composers (higher external complexity) and that has relatively few commodity products (lower internal complexity) may find itself relatively close to the connector quadrant, running the risk of a more chaotic regime where it has to innovate frequently to respond to changes in the platforms of their composers and product improvements of competing complementors. To improve and to move closer to the edge of chaos, the firm may benefit of the combined effects

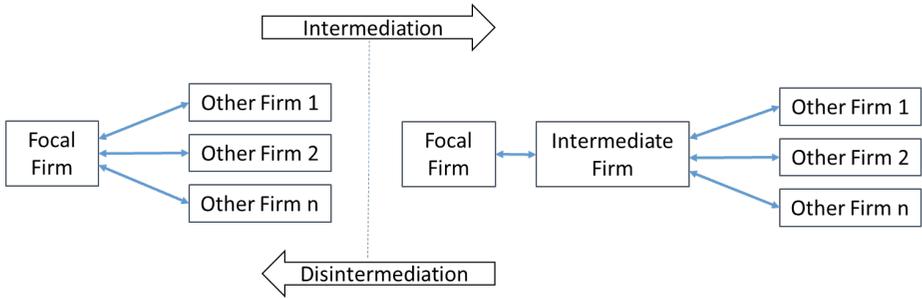


Figure 5-5. Intermediation decreases external complexity and typically to a lesser extent the internal complexity for the focal firm. The reverse happens with disintermediation.

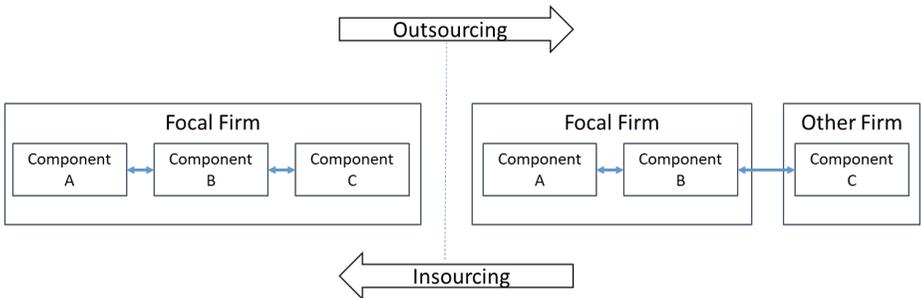


Figure 5-6. Outsourcing increases external complexity and decreases internal complexity. The reverse happens with insourcing.

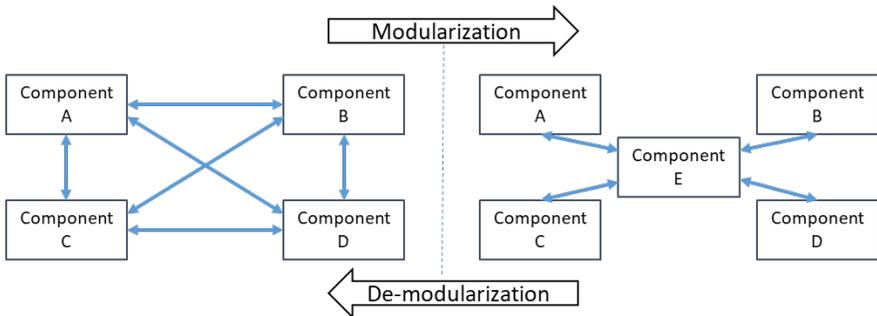


Figure 5-7. Modularization decreases internal / external complexity (depending if the components that are modularized, are inside / outside the focal firm). The reverse happens with de-modularization.

of intermediation and acquisition. The firm may start to work with an intermediary firm that manages the compatibility between its products and the different platforms of composers, decreasing its external complexity. The firm may also acquire another complementor to combine their components and create a more unique product portfolio, increasing their internal complexity.

- D A composer firm that finds itself close to the completer quadrant, with the risk of becoming too rigid in terms of innovation regime, may apply modularization on its internal components to decrease its internal complexity and doing more direct business with more complementors, disintermediating its business model and therefore increasing its external complexity. The effect being that the composer can more easily innovate its products or platform, driven by more direct interaction with more complementors.

Strategic 'Mover type'	Description	Internal complexity	External complexity
Intermediation	Intermediation disconnects the direct links between a firm and its suppliers, customers or other partners in its value chain, by introducing an intermediary firm, decreasing external complexity and to some extent internal complexity by moving some of the firm's internal parts to the intermediary. The opposite happens in case of disintermediation.	-	-
Disintermediation		+	+
Outsourcing	Outsourcing internal parts to another firm decreases internal and increases external complexity. And vice versa for insourcing, acquiring another firm or merging with another firm.	-	+
Insourcing (or Mergers & Acquisitions)		+	-
Modularization (and/or Standardization)	Making a product more modular implies that parts become less interdependent, often through a standardized and common interface between components. If modularized components are internal, then internal complexity will decrease. Likewise, if modularized components are external, then external complexity will decrease. De-modularization or integration will have the reverse effect.	+	+
De-modularization (or Integration)		-	-

Table 5-1. Paired mover types and their general effect on internal and external complexity.

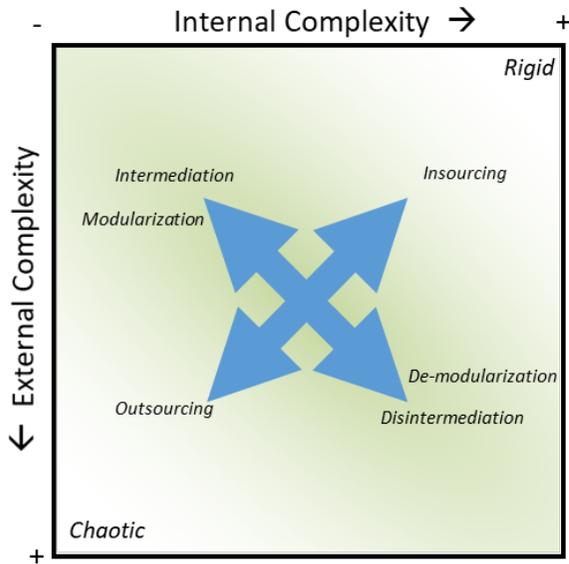


Figure 5-8. Ecosystem complexity model with paired strategic mover types. The arrows indicate the general direction of the move, from any starting position in the ecosystem (not only from the central position), with the exact impact and direction of the move depending on the specifics of a strategic move instance.

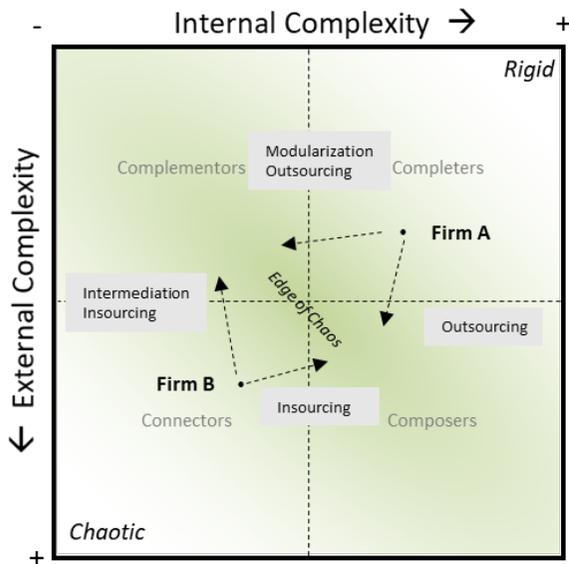


Figure 5-9. Examples of ecosystem position evolution expressed in strategic mover types. Firm A starts as a completer and firm B start as a connector, each moving either towards a composer or a completer position as the ecosystem evolves, each firm applying one or more of the indicated strategies.

Considering the different strategies and types of ‘movers’ that impact complexity and therefore a firm’s ecosystem position, it can be observed that in case of the (de)modularization and (dis)intermediation strategies, the changes in internal and external complexity occur in the same direction: if internal complexity increases then external complexity also increases, and if external complexity decreases then internal complexity also decreases, and *vice versa*. Only the strategy ‘insourcing’ versus ‘outsourcing’ tend to move internal and external complexity in opposite directions. These differences in the direction of impact are highly relevant in achieving a balance between internal and external complexity – required for firms to be as close to the edge of chaos as possible, establishing a regime that provides both stability to exploit and variety to explore. If internal and external complexity are not yet balanced, then the innovation strategy of a firm is in essence the application of one or more of the mover types, to achieve this balance in the envisioned future.

Balancing internal and external complexity brings firms closer to the edge of chaos. This implies that for each firm category, similar movement patterns can be expected. See figure 5-9. Completers and connectors typically move towards either composer or complementor positions. To make this move, firms must apply one or more of the abovementioned strategies. For instance, let’s assume that a completer (firm A in figure 5-9) has managed to achieve a reasonable level of performance. But because of its high internal complexity, the firm finds it hard to keep up with new innovations from competing firms, perhaps from more flexible, newly emerging composers and complementors. Assuming that firm A seeks to change its position towards the composer quadrant, it must manage to increase its external complexity and, depending on its current position, somewhat decrease its internal complexity as well. This could be accomplished, for example, by following a moderate outsourcing strategy, for instance by selling off business units that are suitable for becoming independent complementors, which would also decrease the internal dependencies of the completer.

The completer (firm A in figure 5-9) may also opt to become a complementor, requiring to reduce its internal complexity and also to a lesser extent increase its external complexity. This could be achieved on the one hand by increasing the modularity and the adoption of standards, making it also more feasible to have their components work together with components supplied by other firms. Typically, this move would be accompanied by selling off or ‘outsourcing’ some of its more generic components to composers, or by discontinuing and replacing some of their ‘home made’ own components by components that are supplied by composers.

A firm that starts out as a connector firm (firm B in figure 5-9) may become less relevant once completers start to disappear (or move to other positions, like firm A) and once a few composers start to dominate the ecosystem. The connector may then choose to become a composer itself, typically by changing its offering to a platform in its own right, which could be an autonomous development or more likely a strategy in which the connector is acquired by an existing composer firm, or –less likely– the other way around. The incentive for the existing composer firm to acquire the connector could be to make their platform or infrastructure available to a larger group of complementors, thanks to the interfacing provided by the connector. Alternatively, the connector firm could pursue the strategy of becoming

a complementor, which again would require an extension of its current offering by some more specific product or service, leveraging its current connectivity to one or more composer offerings. The latter is then no longer the primary business of the connector, but rather a head start to create its own complementor offerings. As an alternative the connector could acquire an existing complementor. If a connector is more like a broker, connecting other firms and their services with each other, then it may also choose to focus on a more specific role, by working more locally or by collaborating with other intermediators to reduce its own external complexity, moving towards the complementor quadrant.

Other ecosystem 'journeys' may occur with composers or complementors that seek to move towards the center of the edge of chaos, where performance levels on average are maximized (Kauffman & Johnson, 1991). For composers this implies that they would have to decrease both their internal and external complexity, but not so much to become a complementor. This could be achieved for instance by choosing to replace some of their own components, through modularization and standardization, by externally provided components. However, too much modularization or 'openness' of their offering would push the composer towards a connector position, reducing its performance. Likewise, complementors could try to increase their internal and external complexity, without becoming a composer. For instance by synthesizing their parts to make their product or service more unique with a higher performance level.

It should also be noted that the above mover types are often also 'ecosystem reconfigurators'. They do not only impact the firm's position in the ecosystem, but also impact other actors, by creating new actors (for example in the case of intermediation) or removing existing actors (for example in the case of acquisitions), or by significantly moving parts from one actor to another (for example in the case of outsourcing). When firms employ these mover types, they do not take the existing ecosystem for granted but are actually reconfiguring the ecosystem into a new structure.

Finally, we observe that the abovementioned mover types are neither new nor uniquely purposed for ecosystem moves. In fact, many of these strategies have existed without any notion about ecosystems. Instead, they have been applied from financial, productivity or other perspectives (Elmuti, 2003). For example, many outsourcing initiatives have been driven by the aim to reduce capital or operational expenditures (Kotabe & Mol, 2009). The current analysis demonstrates that these strategies' impact is not limited to such perspectives, and that the impact on a firm's ecosystem position and (co)innovation performance should be an integral part of strategic decision making.

5.6 A Lifecycle Perspective on Ecosystem Positions

While the ecosystem perspective is primarily developed to analyze positions and strategies of individuals firms, ecosystem thinking can also be related to industrial dynamics. In particular, one can ask the question whether the relative presence of different types of ecosystem positions (composer, complementor, connector and completer) is expected to exhibit certain patterns over time.

Among the canonical models of industry evolution is the product lifecycle developed by Abernathy and Utterback (1978) (for a review, see Murmann & Frenken, 2006). This model essentially applies to manufacturing industries and centers on the evolution of a product from an early stage in which many small firms compete with different designs and demand is low, to an intermediate stage in which a dominant design emerges and costs go down and sales go up, to a maturity stage in which few firms dominate the industry and sales level off.

Some propositions can be derived when combining the lifecycle perspective with the ecosystem perspective. First, in the early stage, sales are low and products are not standardized nor modularized. Hence, there is little scope for outsourcing components to suppliers (Klepper, 1997). This would imply that early in the lifecycle of ecosystems or in immature ecosystems, completers should do well, as there are hardly any other firms to do co-innovation with, leaving no other option than doing most innovation internally. In fact, it may be the only viable way in the context of highly complex component assemblies without standards yet. However, in the long run, with sales rising and standardization increasing, completers seem to be bound to lose their position, and to survive they must either become a composer or complementor. In both cases the completer must give up, open up or sell off parts of its offering to other firms. The case of IBM is an example: starting in the 1950s and 1960s in a new and thinly populated computing ecosystem, IBM had basically no choice but becoming a completer. Only after increased competition and the emergence of standards, IBM choose (or had no other choice than) to become a composer, creating a network of complementors to leverage their computing infrastructure and software platforms. An example being the modularization of their 360 system (Baldwin & Clark, 2000). In another example case (Adner, 2012), Nokia was one of the first to introduce a smart phone on the market, with almost all of its hard- and software parts produced by Nokia itself, by lack of other firms that could do this at that moment in time. But even then, in its completer position, Nokia was not able to do all necessary innovations alone or create sufficient momentum to incentivize other firms. Amongst others, Nokia was not in a position to create a 3G mobile network with sufficient bandwidth on its own, for which it depended on national governments and telecommunication companies. Likewise, Nokia was not able on its own to develop a critical mass of content and apps to make their offering attractive enough for consumers. In the end, Nokia had to give up its completer position and, perhaps by a lack of size and capital, reverted back to a role of complementor, selling off parts of its business to Microsoft and other firms.

Connectors may do well in ecosystems in which multiple open or proprietary standards are starting to emerge. This is typical for the intermediate stage in a product lifecycle, but before a single dominant design has emerged (Murmann & Frenken, 2006; Tsai, 2018). The *raison d'être* for connectors will however diminish when a single or few dominant designs emerge and some standards start to dominate while others disappear. Like initial completers, a connector firm must choose to become a complementor or a composer as the intermediate stage of the product lifecycle is ending. A connector firm can do so by either narrowing its offering by connecting to only a single composer, which moves the firm more to a complementor position, or by expanding its offering for others to build on, yet focused on a single standard,

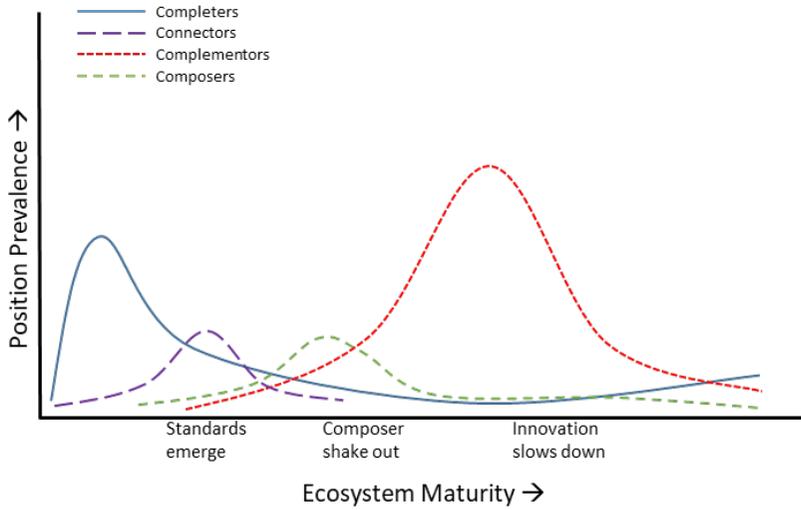


Figure 5-10. Illustrative evolution of ecosystem position prevalence and interaction between positions at different ecosystem maturity stages.

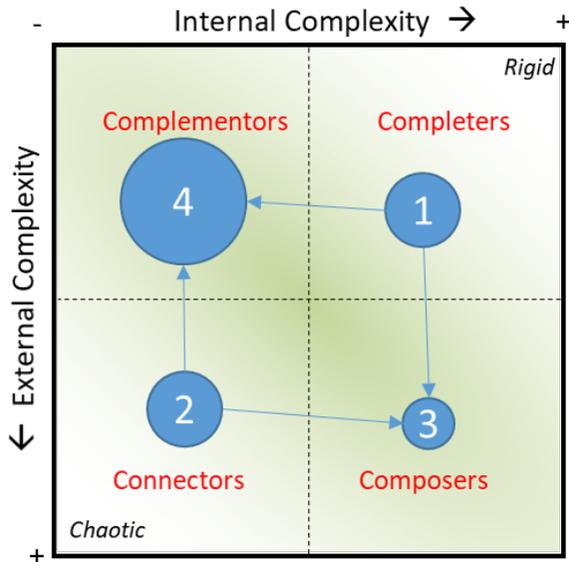


Figure 5-11. Ecosystem evolution: in early stages, companies have no other choice then to start as completer (1). To combine completer solutions, connectors emerge (2). But once composers come into being (3), both completers and connectors cannot sustain their position, and eventually have no other choice then to become either a complementor (4) or also a composer. The size of the circles is a qualitative indication of the number of firms in each category, relative to the other categories.

which moves the firm to a composer position. In the former case, the firm commits to a dominant design introduced by another (composer) firm, while in the latter case, it tries to establish a dominant design itself.

With the emergence of dominant designs, we expect firms to take on more composer positions. Such firms will increasingly compete, leading to a shake out, reducing uncertainty for rising numbers of complementors, which in turn may later consolidate as well as a result of competition. Finally, the rate of innovation starts to slow down and firm entries and exits will be progressively lower. Most firms will reach a stable position characterizing a mature industry. At this late stage, some composers may be tempted by opportunities to expand their scope or improve their profit margin by acquiring or competing with their complementors (Farrell & Katz, 2000; Gawer & Cusumano, 2014a; 2014b), gradually turning themselves into completers.

The resulting industry evolution is depicted in figures 5-10 and 5-11, not implying that each ecosystem evolution instance will indeed follow such a pattern. The purpose here, however, is to illustrate how the prevalence of ecosystem positions may change over time, and how these positions may interact. Elaboration on alternative evolutionary patterns and evaluation against real-world examples are left for future work.

Another remark here on ecosystem dynamics is about the notion that ecosystems as a whole also change. An ecosystem may change as a result of exogenous factors or changes in the collective of ecosystems that we call here the ecosphere. Examples of exogenous factors include government policies, societal or political events, technology shocks, the economic climate or climate change. Changes in the ecosphere also occur because of the emergence of new ecosystems or the decline of existing ecosystems, or because of interactions at ecosystem level. The latter, in contrast to interactions between firms within a single ecosystem, results from interdependencies between ecosystems, creating a higher order complexity.

5.7 Conclusion

The ecosystem complexity model provides an analytical framework to understand the position and roles that firms play in their co-innovation business ecosystem, and how these relate to the positions and roles of other firms. The categorizations of composer, complementor, completer and connector help to characterize the strategic positions of firms, in relation with the rigid, chaotic or edge-of-chaos regimes. The model also provides clues as to how firms might be able to strategically move their position (Baum, 1999), through outsourcing vs. insourcing (Frenken, 2006a; Cela et al., 2018), intermediation vs. disintermediation (Adner, 2012), and modularization vs. de-modularization (Frenken, 2006a; Frenken & Mendritzki, 2012; Henderson & Clark, 1990). Such strategic moves are traditionally theorized by considerations about economy of scale, efficiency, business growth and other matters. Applying the complexity perspective on innovation strategies has revealed that such strategies are also highly relevant in moving the position of firms in their ecosystem, supporting effective innovation.

In addition, much of the existing work in management science can be complemented by applying the ecosystem complexity model. For instance, the work of Adner (2012) on innovation strategy stresses the need for firms not to accept their ecosystem as a given. Firms should actively change the ecosystem position of themselves and, if possible, of other firms, possibly by introducing new firms or reconfiguring the ecosystem in other ways, as far as needed to mitigate co-innovation and adoption chain risks. Application of the ecosystem complexity model reveals that such strategies will vary per position category. It also offers a more tangible and structured toolbox of mover types to achieve required positions or reconfiguration.

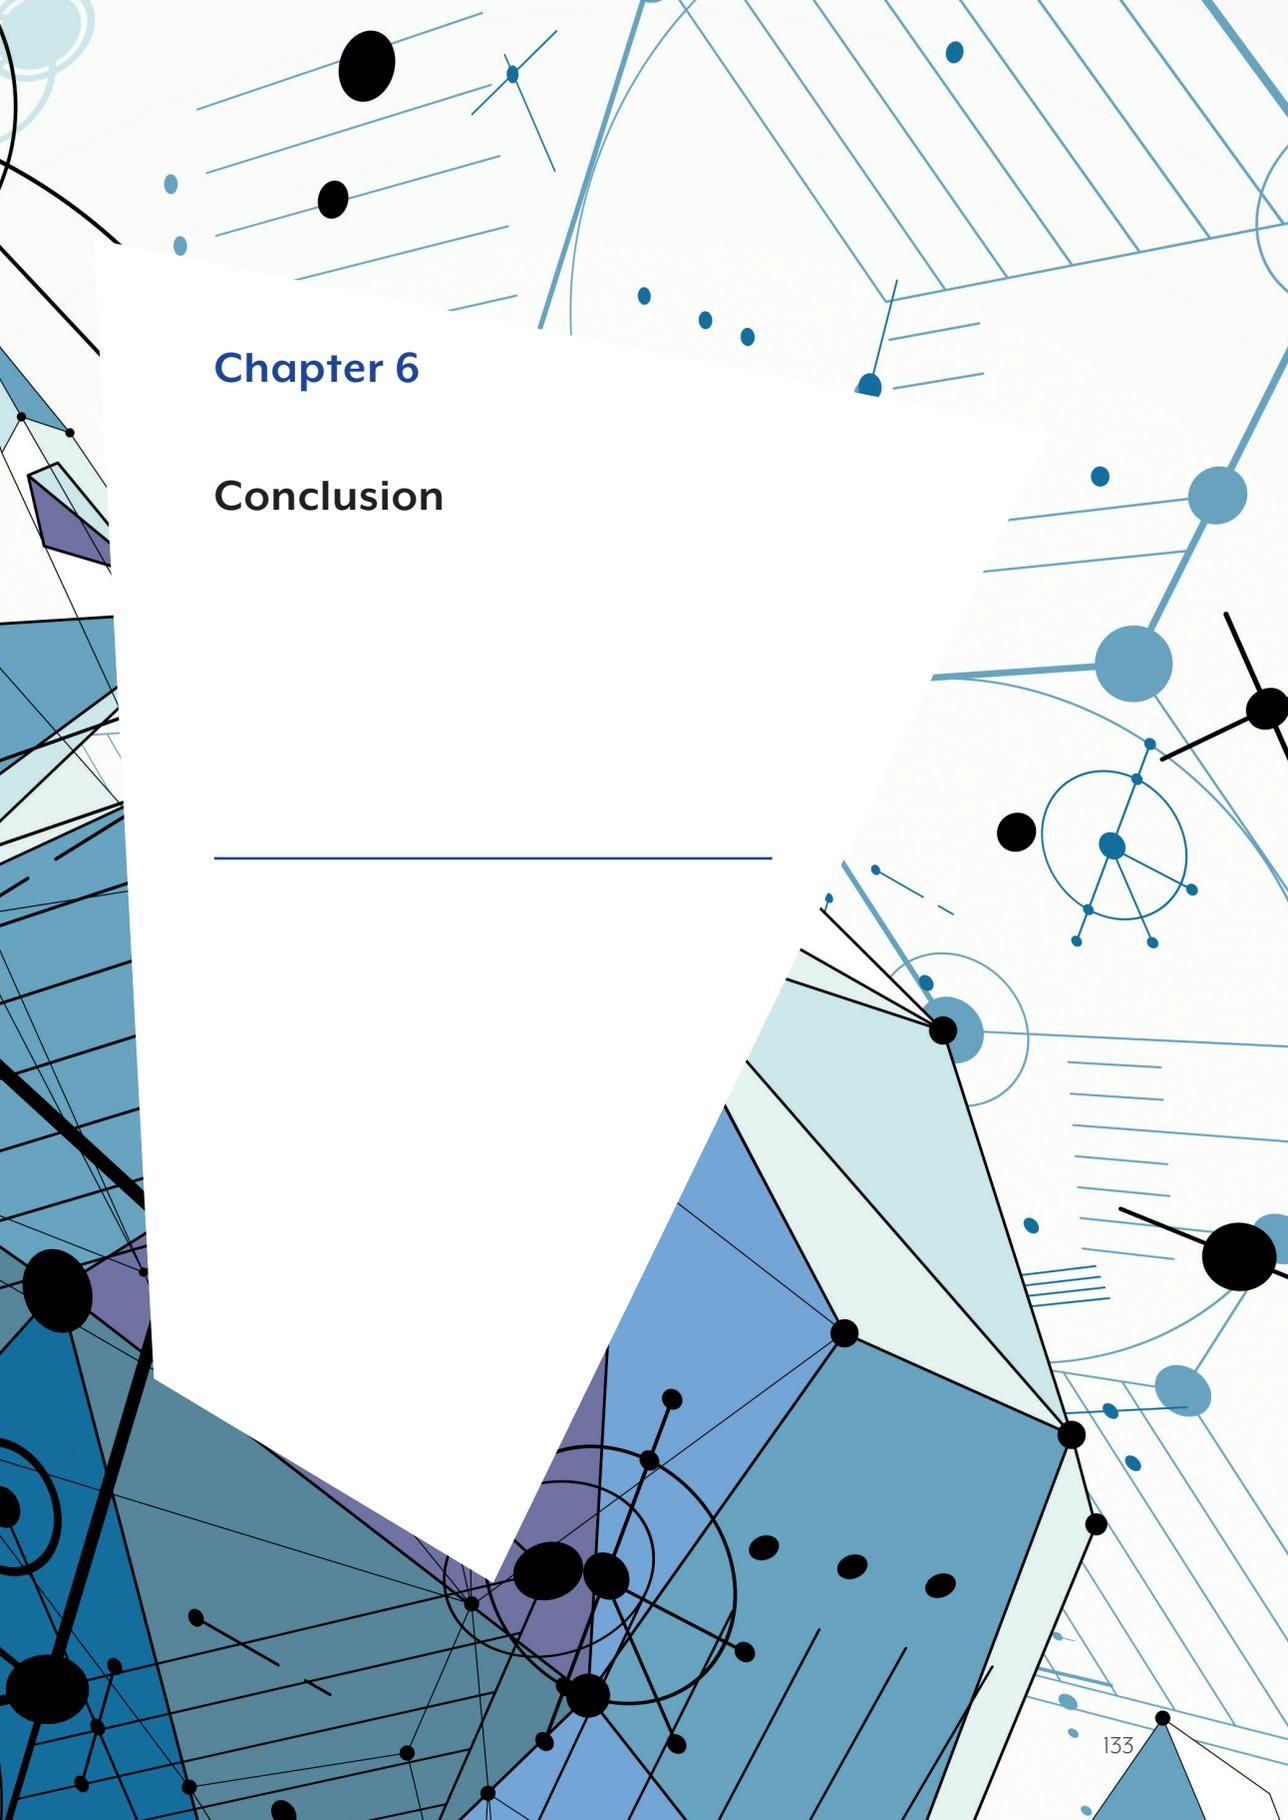
Iansiti & Levien (2004) in their work on innovation strategies, emphasize the necessity of firms to choose and to change their ecosystem role explicitly. Depending on ecosystem conditions, firms may opt for keystone, niche player or dominator roles. As discussed earlier, these roles can now be reformulated and further nuanced in terms of the position categories of the ecosystem complexity model. And, in doing so, the notions of ecosystem regimes, balancing internal and external complexity and applying mover types to change a firm's ecosystem position, are all adding to the depth of insights about effective innovation strategies. *Vice versa*, Iansiti & Levien's work (2004) on ecosystem health and measures on productivity, robustness and variety, which are also addressed in the work of Beinhocker (2007), may be applied to further enrich the ecosystem complexity model. For example, productivity and variety are fostered by composers through their support of complementors, whereas completers decrease variety and connectors increase variety. Robustness is optimal for firms that are closest to the edge of chaos (see also chapter 4), which is mostly populated by symbiotic composer and complementor firms.

The ecosystem complexity perspective further speaks to the product lifecycle model of industry evolution. We proposed that the relative frequency of different types of ecosystem roles (composer, complementor, connector and completer) is expected to change over time from an initial fluid stage with completers and connectors to a more stable configuration of few composers responsible for the dominant design(s) and its standards and many complementor firms to complement the dominant design(s). The reformulation of the product lifecycle model in terms of the ecosystem complexity model does not alter the basic tenets of the product lifecycle model, but rather generalized the model to services industries and platform-based businesses.

In practice, the insights from the current work imply that the complexity and ecosystem perspectives provide a relevant 'wide lens' (Adner, 2012) to both innovation policy makers at government level and innovation managers at the level of individual firms. In essence, the ecosystem perspective underlines the necessity for firms to look beyond their internal operations, realizing that they do not operate or innovate in isolation, but are in fact highly dependent on other firms. Moreover, their ecosystem position is very much dependent on their strategies on changing internal and external complexity levels, which should be balanced for an optimal mix of stability and flexibility, at the edge of chaos. In addition to strategies at firm level, it is also clear that ecosystem themselves are not static. Ecosystems come and go, and evolve during their lifecycle. Moreover, ecosystems also interact with other ecosystems and are of course influenced by external factors like government policies, subsidies, regulations and standardizations.

As the current ecosystem complexity model is largely inspired by work in theoretical biology on ecosystems and co-evolution (Kauffman, 1993) (Kauffman & Johnsen, 1991), amongst others assuming that complexity structures can be taken as random configurations, it is desirable that the propositions in this study are further scrutinized through theoretical simulation studies, and empirical work. In such future work, the current generic and stylized ecosystem complexity model may need to be adapted for specific industry types, like manufacturing or service industries with their idiosyncratic dependency networks between firms and product architectures. Also, the current work is agnostic to possible regional differences and may need to be made more specific to make it practically more useful.





Chapter 6

Conclusion

6.1 Introduction

The common theme in previous chapters was the notion that increasingly, innovation no longer happens in isolation within single firms. Globalization and digitization (Dickens, 2015) and growing technological complexity both enable and require the involvement of multiple firms in innovation. For these and other reasons, management science increasingly recognizes the relevance of business ecosystems (Moore, 1993; Adner 2012; lansiti & Levien, 2004).

The framework of ecosystem perspectives on innovation that was earlier explained in the introductory chapter, see figure 6-1, identifies four perspectives, of which the following three are in scope of this work: contest innovation, market innovation and co-innovation. Contest innovation (Adamczyk et al., 2012; Terwiesch & Xu, 2006) refers to firms that organize innovation competitions, inviting internal units or external units to participate, with all participants addressing the same challenge. Market innovation refers to the situation in which firms compete with similar products or services. Firms may do their own individual innovation, or may attempt to imitate innovations from leading firms (Nelson & Winter, 1982). Co-innovation refers to firms working together, with each firm innovating its own components, which can be combined with other firm’s components into products or services (Kauffman & MacReady, 1995; Kauffman et al., 2000). Each of these perspectives is addressed in chapters 2, 3 and 4 respectively, with chapter 5 providing a more integral analytical framework for innovation strategy in business ecosystems, as indicated in Figure 6-2.

Here the same perspectives will be used to summarize results, to discuss managerial and policy implications, and to provide suggestions for future work. Finally, overseeing the different perspectives together, common and more general conclusions will be drawn.

	Single organization	Multiple organizations
Complementary	Conventional innovation	Co-innovation
Competing	Contest innovation	Market innovation

Figure 6-1. Ecosystem perspectives on innovation

Chapter 5	Single organization	Multiple organizations
Complementary	Conventional innovation	Co-innovation Chapter 4
Competing	Contest innovation Chapter 2	Market innovation Chapter 3

Figure 6-2. Structure of the thesis.

6.2 Contest Innovation

6.2.1 *Summary of results*

The main result of our study on innovation contests (Adamczyk et al., 2012; Terwiesch & Xu, 2006) is the insight that it pays off to have actors who participate in an innovation contest to share knowledge between them. The optimal extent of knowledge sharing must be found between the extremes of on the one hand no sharing at all – leading to wasted efforts and lock-ins (Nooteboom, 2000; Wuyts et al., 2005) – and on the other hand exchanging all available knowledge all the time – leading to a higher risk of premature dominance of a suboptimal approach. The more complicated the problem for which a solution is sought, the bigger the impact knowledge sharing has on performance levels (Hansen, 1999). These results are quite robust for different ways, extents and frequencies of knowledge sharing.

Furthermore, results indicate that innovation contests provide better results than conventional or internal innovation, underpinning the growing popularity of such contests (Terwiesch & Xu, 2006; Adamczyk et al., 2012) and similar applications in crowdsourcing innovation (Franzoni & Saueremann, 2014) and open innovation (Chesbrough, 2003a; 2003b; von Hippel, 2005).

6.2.2 *Managerial & policy implications*

Above results give rise to the need that both innovation managers and policy makers consider to replace or combine classic 'serial' innovation more often with 'parallel approaches' to innovation, either by internal contests or by leveraging the ecosystem of a firm in an external contest. The exception to this are cases in which innovation challenges are too big and too expensive to organize with multiple parallel approaches or cases where such contest would necessitate the disclosure of sensitive knowledge by the organizing firm (Bogers, 2011). In such cases, it is feasible nor desirable to have multiple parties attempting to do the same thing – it is only feasible to combine all available budget and resources to address the challenge together, following a single serial approach.

These exceptions notwithstanding, another implication is that any parallel approach to innovation can benefit from explicit attention to how and how often participants share knowledge. In practice, this implies the application and monitoring of guidelines to regulate or influence the way knowledge sharing takes place. It also implies that firms may organize and facilitate knowledge sharing pro-actively, fostering open innovation, but balanced with the need for variety by keeping multiple approaches alive.

6.2.3 *Limitations and future research*

In this context it was assumed that contest organizers have a clear understanding of the problem to be solved, taking the underlying architecture and complexity as a static given. Clearly, this is limiting the relevancy of findings and applicability of implications. In future work, it would be interesting to see how the model could be extended. For instance by using the NK-model (Kauffman, 1993; Levinthal, 1997) as a starting point, to see how knowledge sharing can be optimized when also dependencies between components or even the structure of these dependencies – the solution architecture – come into play.

It is also recommendable to do further research on the identification and application of guidelines, tools and other means to effectively manage a balanced approach to knowledge sharing in innovation contests and related innovation approaches. Future work may also address questions on the optimal number and size of contestants, certainly if there is a trade-off to be made between those two measures (Levitan et al., 2002). In other words, would it be more effective to have just a few, bigger units doing innovation, but with less other contestants to obtain knowledge about alternatives approaches, or would it be more effective to have more, smaller units with more alternative approaches available, but with less resources to exploit each unit's own approach?

6.3 Market innovation

6.3.1 *Summary of results*

The main theoretical implication of our work from a market innovation perspective, concerns the key role of social networks among firms in supporting effective imitation, 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 dyad level of two firms, we further have been able to integrate the theory of small-worlds in collective learning and the proximity theory regarding learning between any two firms, finding robust theoretical support both for the thesis of optimal cognitive proximity and for the substituting role of social proximity in bridging cognitive distance.

Introducing the notion of internal complexity (Kauffman, 1993) has revealed that imitation (Rivkin, 2000) between competing innovators is most relevant for highest complexity, unnecessary for lowest complexity and most rewarding at intermediate complexity levels. If complexity is low, then a firm can do its own innovation, without the need to imitate from competitors. If complexity is high, a firm can quickly get stuck in a suboptimal position, making imitation from others an important tool to find another more promising avenue. However, if complexity is high, so is the likelihood of error making during the imitation process, as small mistakes may cause large performance degradations. At intermediate complexity, imitation is both more feasible and therefore also more rewarding in terms of performance gain compared to high complexity, and more relevant compared to lower complexity.

Informal knowledge sharing between employees from competing firms who are acquainted, can be an important channel for imitation efforts to be effective (Breschi & Lissoni, 2009). We have modeled this by a social network that defined such informal knowledge relations and assuming that an imitating firm that is closer in the social network to the imitated firm, makes fewer errors than when imitating a firm that is more distant in the social network (Boschma, 2005). From the simulation results, however, it was found that imitation at intermediate social proximity offers optimal performance gains, at least for intermediate and higher complexity. In addition, the same was found, but less prominent, for the cognitive proximity between firms. In both cases of proximity, too much proximity seems to generate insufficient newness to escape a current suboptimal and unpromising position, as firms at close proximity tend to become ever more similar in their thinking and innovation approach. More

distance is required to find a new approach, although too much distance renders imitation infeasible (Wuyts et al., 2005; Nooteboom et al., 2007; Gilsing et al., 2008; Fitjar et al., 2016).

In addition, when comparing different structures of the network of social connections between firms (Watts & Strogatz, 1998), being the main conduits of imitation, it was found that small-world networks offer best performance. In such networks, the benefit of close social proximity through shortcuts between clusters is combined with redundant pathways within clusters to facilitate knowledge transfer (Schilling & Phelps, 2007). As such, small-world networks outperform both random networks, which have high proximity but lack clustering, and regular networks, which are highly clustered but offer less proximity on average. However, the benefits of small-world networks only become apparent at later innovation stages when performance is already close to peak level, and in case of higher complexity levels. It is only at such levels that error making becomes more costly, making the benefits of clustering more relevant. At lower levels and during earlier stages, random networks offer the quickest way to improve performance.

6.3.2 Managerial & policy implications

As complexity is highly relevant in determining optimal innovation conditions, innovation policy makers and managers should try to establish at least a general sense of whether complexity levels are low versus intermediate or higher. If not low, social and cognitive proximity are important criteria to apply in what should be the explicit identification and selection of competing innovators for imitation. Competitors that are very proximate are less likely to offer new insights, while too little proximity will hamper effective imitation. Also, to compensate for higher cognitive distance and to support imitation, managers may actively identify socially proximate competitors, or seek ways to reduce social distance between their own and competitor employees.

Moreover, if complexity is high and performance levels are high – which especially can be expected in more mature and high-tech industries – network structure matters. Although firms will not be able to fully control the structure of their social network, policy makers on the other hand may try to prevent or reduce too much clustering by explicitly promoting cross-links between clusters. The result of which should be the enablement of the formation of networks with small-world characteristics, while avoiding too much cross-linking that would foster a random structure, diminishing clusters. A typical example of a possible tool to influence network structure formation would be the requirement in innovation subsidy schemes to have project consortia in which two or more firms from different clusters are involved, with their project collaboration creating a cross-link.

6.3.3 Limitations and future research

In the current research, internal complexity is assumed to remain constant. In reality however, the architecture of solutions may change over time (Altenberg, 1994), for instance due to modularization and standardization, in effect lowering the dependencies between constituent parts (Baldwin & Clark, 2000; Fleming & Sorenson, 2001a). Also, over time, dominant designs may emerge (Abernathy & Utterback, 1978; Frenken & Mendritzki, 2012), implying that groups of firms adopt a similar

architecture. These complexity dynamics may have an impact on innovation policies and managerial guidelines, which may need to change along with these dynamics, during an industry's lifecycle. For instance, it is worth to investigate if random networks offer preferable conditions once an industry has adopted significant standardization or if clustering may lead to multiple coexisting dominant designs.

Furthermore, the structure of the social network (between employees of different firms) is considered to be exogeneous, in the current work. In reality, social networks may be influenced more purposefully, for example through the formation of R&D alliances (Cowan & Jonard, 2009). Future work is welcome to further clarify and identify ways in which firms can actively influence their social network, instead of considering the network as a given. Options may include the active 'poaching' of competitor employees, or increasing geographic proximity with competing firms to facilitate social network building between employees of different firms.

In addition, it was assumed that firms only do incremental innovation by local search and only do radical innovation by imitation (Nelson & Winter, 1982; Lieberman & Asaba, 2006). In the real world, firms are likely to apply other approaches too (for a review, see Ordanini et al., 2008), such as initiating their own R&D initiatives or participating in R&D alliances (Cowan & Jonard, 2009) for radical innovation. Future work may enrich the current model to also account for such other innovation approaches. Such as by incorporating exploration through internal search as in the original NK-model (Kauffman, 1993; Levinthal, 1997). Or by taking into account that doing innovation within alliances is arguably quite distinct from imitation. Alliances are cooperative structures aiming to generate new knowledge for both partners through recombination ('crossover') rather than a new solution to the imitating firm only (Cowan & Jonard, 2009). A future model may systematically compare the outcomes of exploration via imitation versus exploration via alliances, to investigate the generalizability of our results. Such a model may also address the question of 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 benefit more from imitating others.

Another possible future model extension is to incorporate cost. In our model, we abstracted 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). Still another extension concerns the investigation of environmental turbulence on exploration and exploitation, which can be integrated with the NK-model by making fitness level noisy (Uotila, 2017). Here, a key question holds whether imitation, then, is still as effective as a means to conduct exploration, as fitness levels of competitors then only convey noisy information.

Finally, a promising avenue for future work lies in considering more, multiple proximity types and to see how they impact knowledge transfer in the context of innovation. In addition to social and cognitive proximities, other proximities worth considering include geographical, organizational and institutional proximities (Boschma, 2005). None of these proximities is likely to be sufficient on its own to explain how imitation and knowledge transfer take place in reality, but by considering these proximities

together, possibly by applying a multi-layered network model (Bródka et al., 2012), a further and deeper understanding may be gained.

6.4 Co-Innovation

6.4.1 *Summary of results*

From earlier work we know already that firms involved in co-innovation in their business ecosystem may find themselves to operate in one of a number of quite different regimes, just like organisms in their co-evolution in biological ecosystems (Kauffman & Johnsen, 1991). These regimes are either chaotic, with continuous change occurring, making sustainable progress very hard to achieve, or rigid, with firms having mostly internal and much less external dependencies, resulting in inertia and limited change. In between is the 'edge of chaos' regime, in which flexibility is balanced with stability. In our work here, further underpinning is given about the fact that performance levels are indeed highest in the edge-of-chaos regime, compared to lower performance in the other regimes.

The main theoretical implication of our work holds that the conditions that give rise to the edge-of-chaos regime, do not only depend on the ratio between internal and external complexity (i.e., K versus C), but are also influenced by the network structure of interdependencies between firms in the ecosystem. In summary, it is observed that regular, highly clustered networks move the ecosystem to a more rigid regime, whereas more random network move the ecosystem to a more chaotic regime. Small-world networks, combining characteristics of both regular and random networks, seem to provide a structure that is best suited for the edge of chaos regime. In an extension of the model, we further showed that if firms follow simple networking rules, the whole ecosystem tend to evolve endogenously towards a small-world network structure, providing a further theoretical understanding why small-world network structures are empirically often observed.

In addition, it was found that edge of chaos regimes not only provide optimal performance, but also best resilience to shocks that are applied to the ecosystem. In edge of chaos regimes, shock impacts are limited compared to more chaotic regimes. Moreover, these impacts provide an opportunity to escape from local sub-optima, allowing firms to repair or even find higher levels of performance, in contrast to more rigid regimes. Interestingly, the distribution of the size frequency of shock impacts is closest to a power law in edge of chaos regimes. Apparently, the occurrence of a power law distribution is a reflection of the ecosystem being in a state of self-organised criticality (Bak et al., 1998) in which stability and adaptability are perfectly balanced.

6.4.2 *Managerial & policy implications*

A key implication for both innovation managers and policy makers is the finding that ecosystems and the position of individual firms therein, are highly relevant. At firm level, managers should be aware of how internal and external complexity influence their own capability to co-innovate and that of their innovation partners, and determine in which ecosystem regime they operate. Explicitly managing internal and

external complexity levels (Baum, 1999), for instance by changing the architecture of dependencies between components in a product or by in- or outsourcing of components, allow a firm and its partners to move towards the optimal edge of chaos regime.

At ecosystem level, policy makers and large, dominant firms should be aware how the network structure that emerges from the interdependencies between firms, also is a significant factor influencing the ecosystem regime. Encouraging firms towards more clustering, pushes the ecosystem in the direction of a more rigid or less chaotic regime, and *vice versa*. Policy makers can monitor and assess ecosystem regimes by keeping track of complexity levels and the occurrence of innovation cascades, including their size and size distribution. This allows to exert influence in the right direction, to get closer to the edge of chaos, which in turn should result in both higher performance and higher resiliency levels.

Nevertheless, the ability to explicitly manage or even influence the ecosystem regime from a top-down perspective or at level of an individual firm, should not be overestimated, as most ecosystems and their regime typically emerge spontaneously or at least are not strongly orchestrated by a single actor, at least not during the initial lifecycle stages of the ecosystem. This may change once an ecosystem gets more dominated by one or a few larger firms.

6.4.3 Limitations and future research

The current simulation studies are capable of reproducing and relating multiple stylized facts about clustered innovations and their distribution in size and time (Arthur, 2007; Verspagen & Duysters, 2004; Dosi, 1988). Results are also in line with empirical findings on the ubiquity of small-world inter-firm network structures (Marion, 1999; Ter Wal, 2013; Chen & Guan, 2010). Nevertheless, the confidence in our results should be strengthened further in empirical work that is more explicit and complete in addressing aspects that seem to be most relevant. Preferably, empirical work should try to gauge both internal and external complexity of ecosystem firms and their performance, determine the structure of the ecosystem network and identify time-clustered innovations, innovation cascades and their size.

Another promising area for future work lies in the fact that ecosystems are dynamic: new firms are born, and other firms go bankrupt. These dynamics affect the structure of the ecosystem network, by creating and changing the connections or interdependencies between firms. It would be interesting to see how current models can be extended to represent the dynamics of firm life cycles, or other events and (evolving) characteristics at firm level, such as alliance or network collaboration strategies or more implicit preferences, and to see how these may influence the emergence of regimes or other characteristics at ecosystem level.

6.5 Combining perspectives: the ecosystem complexity model

The three ecosystem perspectives on innovation – contest innovation, market innovation and co-innovation – do not rule each other out. In the real world, it is quite common that firms that share an ecosystem are simultaneously co-innovating

with some firms, competing with other firms and perhaps are also involved in contest innovations with still other firms. It may even occur, especially in larger conglomerate firms with multiple divisions, that one division is co-innovating with a different firm, while another division is competing with that same firm.

6.5.1 Summary of results

To bring the ecosystem perspectives together, the 'ecosystem complexity model' was introduced in chapter 5, using internal and external complexity as main dimensions (Kauffman & Johnsen, 1991). Using these dimensions, the position of a firm in its ecosystem can be categorized as completer (high internal, low external complexity), connector (low internal, high external complexity), composer (high internal, high external complexity) or complementor (low internal, low external complexity). In addition, ecosystem regimes can be characterized by using the same dimensions. High internal complexity and low external complexity yield a rigid regime. Low internal and high external complexity yield a chaotic regime. The edge of chaos regime is characterized by a balanced internal and external complexity (Kauffman & Johnsen, 1991).

The main theoretical contribution of the current work concerns the analytical underpinning of the emerging work on business ecosystems in general, and innovation ecosystems in particular (Adner, 2012; Iansiti & Levien, 2004), by abstracting on the notions of ecosystem regimes and complexity levels that were introduced in seminal work on the NKCS model by Kauffman and Johnsen (1991). In addition, by applying the complexity perspective we were able to extend the formative product lifecycle model (Abernathy & Utterback, 1978) from being mostly applicable to manufacturing industries to also other industries, including service oriented industries and platform oriented industries.

The ecosystem complexity model has proven to be useful in a number of ways. First, individual firms can use the firm categorization to gauge not only their own position, but also the position of other co-innovating firms. This reveals typical co-innovation partnerships, such as between a connector and multiple completers, with the connector providing a product or service, such as an adapter or convertor, that brings otherwise non-compatible completer solutions together, despite a lack of standards. Other typical co-innovation partnerships are those between a composer and multiple complementors, with the latter firms offering value adding solutions on top of a platform or infrastructure that is provided by a composer, typically using standards to facilitate their interfacing.

Second, the ecosystem complexity model offers insights about how firms can and should develop their position, keeping away from too rigid or too chaotic regimes, which is especially a risk for completers respectively connectors, and moving towards the edge of chaos, choosing a role as either composer or complementor. The model also helps in identifying strategies that firms may utilise in moving their position, in all cases changing their internal or external complexity, or both (Baum, 1999). Such strategies include insourcing (or mergers & acquisitions) and outsourcing of the components of their offerings, (dis)intermediation between a firm and other firms, and changing the architecture of their offering by modularization or de-modularization (or integration).

Finally, the ecosystem complexity model helps to identify and understand the dynamics of firm positions and how ecosystems are likely to evolve in time, in essence extending the product lifecycle model developed by Abernathy and Utterback (1978). The optimal position of a firm seems to vary with the maturity of the ecosystem. During early immature stages, firms basically have no other choice than to be a completer, by lack of other firms to co-innovate with. Once more firms have entered the ecosystem, connector firms become relevant to make solutions work together, until standards and platform products and services, provided by a number of emerging composer firms, diminish the relevancy of connector services and facilitate the development of complementor firms.

6.5.2 Managerial & policy implications

At the level of individual firms, innovation is likely to benefit from a more ecosystem aware innovation strategy, with firms explicitly assessing their current position and moving to a desired position, typically as close to the edge of chaos as possible. Such strategies can be supported by activities that affect a firm's internal complexity and external complexity (Baum, 1999). These activities include insourcing/outourcing, (dis)intermediation and (de)modularization.

However, to manage innovation effectiveness, firms should also be aware of the position and development of other complementary firms. Notably, composer firms can only be successful in the presence of complementor firms, and *vice versa*. Also, given the inherent interdependencies between firms in a co-innovation ecosystem, if most firms operate in a chaotic regime it is less likely that a single firm or even a small cluster of firms is able to move to a more rigid regime, or *vice versa*. In other words, firms must realize that their individual capability to move their own or other firms' positions, is likely to be limited, unless firms manage to coordinate their movements either collectively or through orchestration by typically a large, dominant firm.

At innovation policy level, the ecosystem perspective on co-innovation contributes to a deeper and richer understanding of innovation effectiveness and dynamics. Using these insights, policy makers can try to influence the ecosystem as a whole, by explicitly promoting or discouraging certain regimes. For example, an ecosystem that has many completer firms is likely to be too rigid, with slow innovation and even inertia, resulting in degrading performance. This may be changed by policies or laws and regulations that help to break down such completer firms into smaller firms, opening up the way for composer and complementor firms that are closer to the edge of chaos. Or, as another example, an ecosystem that has many interdependent firms – with high external complexity – may create a chaotic regime in which firms fail to book sustainable progress in their innovation because of ever changing conditions, and in which resilience to ecosystem shocks is also poor. The latter is not only relevant from an economical or performance perspective, but also from a societal perspective, as too many disturbances and bankruptcies often result in unemployment or other negative effects on welfare. To move the ecosystem more towards a less chaotic regime, back to the edge of chaos, policy makers may for instance actively propagate the development and adoption of one common standard, reducing interdependencies. Alternatively, policy makers can try to stimulate the clustering of firms, through coordination or subsidy schemes, changing the ecosystem network to a more regular structure and pushing the ecosystem towards a more rigid state.

6.5.3 Limitations and future research

The framework we presented in chapter 5 is generic and stylized. The current qualitative framework may be strengthened and further developed through formal modeling and simulation studies. Moreover, the model is agnostic regarding differences between specific industries, markets or countries. One may argue that the current model, with its emphasis on composers and complementors, amongst others, may be particularly relevant to platform-oriented sectors, such as the IT industry. As such, it may be less relevant to industries in which products are less modularized, with innovation mostly taking place within single firms, like completers. It would be interesting to see if the model requires adaptation to account for such possible idiosyncratic industry differences that may emerge from empirical work, or that the model can be maintained as is, fulfilling the expectation that complexity and ecosystem regimes are fundamental to all industries.

Most extant work on innovation in ecosystems has focused on case studies. The current work contributes by abstracting and considering the field from an analytical perspective, relating to a number of case studies for illustration. However, in future work we would welcome a more comprehensive literature review of empirical work, actively using the current framework for classification and other analytical purposes, and evaluating its applicability and usefulness.

6.6 Overall conclusions

In general, the results of this work contribute to theorizing in evolutionary economics, innovation studies, and management science. By applying and combining seminal models such as the small-world networks (Watts & Strogatz, 1998), the NK and NKCS models (Kauffman, 1993; Kauffman & Johnsen, 1991), the product lifecycle model (Abernathy and Utterback, 1978) and by adding the ecosystem complexity model, the current work has resulted in findings that are in line with multiple stylized facts (Arthur, 2007; Verspagen & Duysters, 2004; Dosi, 1988). And although not comprehensively addressed, findings also seem to be in line with existing empirical findings (Ozman, 2009; Marion, 1999; Colovic & Cartier, 2007; Ter Wal, 2013; Chen & Guan, 2010; Cowan et al., 2006; Schilling & Phelps, 2005; Baum et al., 2003; Rowley et al., 2000; Sorenson et al., 2006; Mowery et al., 1998; Gubbins & Dooley, 2014; Adner & Kapoor, 2010; Kapoor & Agarwal, 2017). In a more general scientific sense, our results have demonstrated the value of applying the perspectives of complexity, evolution and ecosystems not only to understanding biological systems, but also to understanding human societal systems and their dynamics, including innovation (Anderson, 1999; Beinhocker, 2007).

Taken together, the identified managerial and policy implications offer a significant extension of the 'toolbox' that managers and policy makers can apply to improve the effectiveness of innovation. These instruments include the guidelines for knowledge sharing between innovation actors in a collective setting, the balancing of social and cognitive proximity in knowledge transfer from competing innovators, the awareness and possible influencing of a firm's network position and overall ecosystem network structure, and the strategic relevance of firm positions in the ecosystem complexity model. In general, it has become only more evident that ecosystems matter to

innovation, requiring explicit managerial practices and policies. Having said this, the abstract nature of our model would necessitate further research to render the models more specific to particular firm or industry contexts as to derive more specific managerial implications.

Traditionally, many innovation practices and policies have been driven by on the one hand internal, single firm considerations or on the other hand macro-economic considerations. The insights that we have derived from complexity theory offer a perspective that is in between these two extremes, at ecosystem level. A perspective that is much needed and deserves explicit attention. Firms that ignore their ecosystem and focus only on their internal innovation processes, are very likely to fail, as demonstrated by a growing number of empirical case studies (Adner, 2012; Iansiti & Levien, 2004). Treating in- or outsourcing decisions, for example, purely from a financial perspective, fails to recognize the impact on complexity, which is both a risk and a missed opportunity for the firm at hand. Not to mention the potentially damaging and cascading effects on other firms in the ecosystem.

Innovation policy makers that consider the economy or market only from a general macro level, treating all firms and ecosystems equally, are likely to be ineffective, as they fail to recognize the differences between ecosystems in terms of their lifecycle, network structure, complexity levels or regime. Instead they should differentiate their policies by contextualizing them for each different ecosystem. Moreover, they should recognize that interdependencies between firms, which are only increasing as a result of digital technology and other trends, are at the very heart of what makes an ecosystem work, or not. Firms come and go, but an ecosystem is more than the sum of its constituent firms – instead it can be seen as an entity that emerges from the interdependencies between its members. An anonymous yet robust entity which existence does not depend on the birth or death of individual firms. An entity that exists at intermediate level between markets and individual firms, whose health and characteristics are highly relevant to innovation performance and other economic and societal goals, and therefore deserves explicit management and policies.

Overseeing all results, the theme of 'balance' stands out as a common denominator. Balance between too little and too much knowledge sharing in innovation contests. Balance between too close or too distant social and cognitive proximity in imitation efforts between competing firms. And, for business ecosystems, the balance between too low and too high complexity levels of interdependencies, between internal and external complexity, and between too much clustering and too much randomness in the ecosystem's network structure.

These conclusions underline the key notion that organizations do not operate in isolation, but can only exist as part of their ecosystem. One may even say that their ecosystem not only facilitates their operations, but actually shapes and defines their very existence. Noting this may seem trivial, but formally theorizing ecosystems has been shown to deliver multiple new insights. What is more, both in science and practice of innovation management and policy making, the ecosystem notion is relatively new or often considered only implicitly.

A deeper understanding of underlying structure, mechanisms and factors is required to understand and – where possible – influence the innovation performance of firms and the health of their ecosystems, to identify and find optimal balances

in sharing, proximity, network structure, complexity and regime. Every actor is part of a network, an ecosystem, without which we cannot thrive in business, innovation and life in general. It therefore only seems to make sense that theories on complexity, evolution and ecosystems, all having much of their roots in the biological realm of nature, have proven to be so fertile in furthering our insights. And likewise, it only seems to make sense that economic science and management science must meet in the middle, between market and individual firm, merging into what may become the science of ecosystems.





References

A

- Abernathy, W. J., & Utterback, J. M. (1978). Patterns of industrial innovation. *Technology Review*, 80(7), 40-47.
- Adamczyk, S., Bullinger, A. C., & Möslin, K. M. (2012). Innovation contests: A review, classification and outlook. *Creativity and Innovation Management*, 21(4), 335-360.
- Adner, R. (2006). Match your innovation strategy to your innovation ecosystem. *Harvard Business Review*, 84(4), 9-107+148.
- Adner, R. (2012). *The wide lens - a new strategy for innovation*. London: Penguin Books.
- Adner, R. (2017). Ecosystem as structure: An actionable construct for strategy. *Journal of Management*, 43(1), 39-58.
- Adner, R., & Kapoor, R. (2010). Value creation in innovation ecosystems: How the structure of technological interdependence affects firm performance in new technology generations. *Strategic Management Journal*, 31(5), 306-333.
- Ahouse, J., Bruderer, E., Gelover-Santiago, A., Konno, N., Lazer, D., & Veretnik, S. (1991). Reality kisses the neck of speculation: A report from the NKC workgroup. *Lecture in complex systems, SFI studies in the science of complexity 4* (pp. 331-353). Reading, MA: Addison Wesley.
- Ahuja, G. (2000). Collaboration networks, structural holes, and innovation: A longitudinal study. *Administrative Science Quarterly*, 45(3), 425-455.
- Akrich, M., Callon, M., Latour, B., & Monaghan, A. (2002). The key to success in innovation part I: The art of interressement. *International Journal of Innovation Management*, 6(02), 187-206.
- Albrecht, H. (1997). *Laserforschung in Deutschland 1960-1970*. Unpublished Habilitation Thesis, University of Stuttgart.
- Altenberg, L. (1994). Evolving better representations through selective genome growth. *Proceedings of the IEEE world congress on computational intelligence* (pp. 182-187) Piscataway NJ: IEEE.
- Altenberg, L. (1997). NK fitness landscapes. Section B2.7.2 in T. Back, D. Fogel & Z. Michalewicz (Eds.), *The handbook of evolutionary computation*. Oxford University Press.
- Anderson, P. (1999). Perspective: Complexity theory and organization science. *Organization Science*, 10(3), 216-232.
- Armstrong, M., & Wright, J. (2007). Two-sided markets, competitive bottlenecks and exclusive contracts. *Economic Theory*, 32(2), 353-380.
- Arthur, W. B. (2007). The structure of invention. *Research Policy*, 36(2), 274-287.
- Ashby, W. R. (1991). Requisite variety and its implications for the control of complex systems. *Facets of systems science* (pp. 405-417). Boston, MA: Springer.
- Axelrod, R., Mitchell, W., Thomas E., R., Bennett S., D., & Bruderer, E. (1995). Coalition formation in standard-setting alliances. *Management Science*, 41(9), 1493-1508.

B

- Bak, P., Tang, C., & Wiesenfeld, K. (1988). Self-organized criticality. *Physical Review A*, 38(1), 364.
- Baldwin, C. Y., & Clark, K. B. (2000). *Design rules. Volume I: the power of modularity*. Cambridge, MA: MIT Press.
- Barnett, M. L. (2006). The keystone advantage: What the new dynamics of business ecosystems mean for strategy, innovation, and sustainability. *Academy of Management Perspectives*, 20(2), 88-90.
- Baum, J. (1999). Whole-part coevolutionary competition in organizations. *Variations in organization science: In honour of Donald T. Campbell* (pp. 113-135). Thousand Oaks, CA: Sage.
- Baum, J. A., Shipilov, A. V., & Rowley, T. J. (2003). Where do small worlds come from? *Industrial and Corporate Change*, 12(4), 697-725.
- Baumann, O., Schmidt, J., & Stieglitz, N. (2019). Effective search in rugged performance landscapes: A review and outlook. *Journal of Management*, 45(1), 285-318.
- Beinhocker, E. D. (2007). *The origin of wealth: The radical remaking of economics and what it means for business and society*. Cambridge MA: Harvard Business Press.
- Bekkers, R., Duysters, G., & Verspagen, B. (2002). Intellectual property rights, strategic technology agreements and market structure: The case of GSM. *Research Policy*, 31(7), 1141-1161.
- Benner, M. J., & Tushman, M. L. (2003). Exploitation, exploration, and process management: The productivity dilemma revisited. *Academy of Management Review*, 28(2), 238-256.
- Berends, H., Deken, F., & Hummel, J. T. (2018). This is NOT VHS vs. betamax! the development of platform strategies in home automation. Paper presented at the *Academy of Management Specialized Conference: Big Data and Managing in a Digital Economy*.

- Birchenhall, C., Kastrinos, N., & Metcalfe, S. (1997). Genetic algorithms in evolutionary modelling. *Journal of Evolutionary Economics*, 7(4), 375-393.
- Bogers, M. (2011). The open innovation paradox: Knowledge sharing and protection in R&D collaborations. *European Journal of Innovation Management*, 14(1), 93-117.
- Boschma, R. (2005). Proximity and innovation: A critical assessment. *Regional Studies*, (39:1), 61-74.
- Boudreau, K. (2010). Open platform strategies and innovation: Granting access vs. devolving control. *Management Science*, 56(10), 1849-1872.
- Bouty, I. (2000). Interpersonal and interaction influences on informal resource exchanges between R&D researchers across organizational boundaries. *Academy of Management Journal*, 43(1), 50-65.
- Brandenburger, A. M., & Nalebuff, B. J. (1997). The added-value theory of business. *Strategy and Business*, 9(4), 4-6.
- Breschi, S., & Lissoni, F. (2009). Mobility of skilled workers and co-invention networks: An anatomy of localized knowledge flows. *Journal of Economic Geography*, 9(4), 439-468.
- Bródka, P., Kazienko, P., Musiał, K., & Skibicki, K. (2012). Analysis of neighbourhoods in multi-layered dynamic social networks. *International Journal of Computational Intelligence Systems*, 5(3), 582-596.
- Bron, C., & Kerbosch, J. (1973). Algorithm 457: Finding all cliques of an undirected graph. *Acm*, 16(9), 575-577.
- Brusoni, S., & Prencipe, A. (2013). The organization of innovation in ecosystems: Problem framing, problem solving, and patterns of coupling. *Collaboration and competition in business ecosystems* (pp. 167-194), Emerald Group Publishing Limited.
- Burt, R. S. (1992). *Structural holes: The social structure of competition*. Cambridge, MA: Harvard University Press.
- Burton, B., Blossch, M., Natis, Y., Malinverno, P., Brand, S., Osmond, N., . . . Valdes, R. (2017). *Capitalizing on your business ecosystems economy*. Stamford: Gartner. Retrieved from <https://www.gartner.com/document/3759164>

C

- Caminati, M. (1999). *Complementary, opportunity and the co-evolution of techniques*. Manuscript, Università di Siena.
- Cantu-Paz, E. (1995). *A summary of research on parallel genetic algorithms*. Urbana-Champaign: University of Illinois.
- Capaldo, A. (2007). Network structure and innovation: The leveraging of a dual network as a distinctive relational capability. *Strategic Management Journal*, 28(6), 585-608.
- Cassi, L., & Plunket, A. (2015). Research collaboration in co-inventor networks: Combining closure, bridging and proximities. *Regional Studies*, 49(6), 936-954.
- Celo, S., Nebus, J., & Wang, I. K. (2018). The role of internal and external complexity in global factory performance: An NKC application. *Journal of International Management*, 24(1), 65-83.
- Cennamo, C., & Santalo, J. (2013). Platform competition: Strategic trade-offs in platform markets. *Strategic Management Journal*, 34(11), 1331-1350.
- Chang, M., & Harrington, J. (2000). Centralization vs. decentralization in a multi-unit organization: A computational model of a retail chain as a multi-agent adaptive system. *Management Science*, 46(11), 1427-1440.
- Chang, M., & Harrington, J. (2003). Multimarket competition, consumer search, and the organizational structure of multiunit firms. *Management Science*, 49(4), 541-552.
- Chao, R. O., & Kavadias, S. (2008). A theoretical framework for managing the new product development portfolio: When and how to use strategic buckets. *Management Science*, 54(5), 907-921.
- Chen, Z., & Guan, J. (2010). The impact of small world on innovation: An empirical study of 16 countries. *Journal of Informetrics*, (4), 97-106.
- Chesbrough, H. W. (2003a). The era of open innovation. *MIT Sloan Management Review*, 44(3), 35-41.
- Chesbrough, H. W. (2003b). *Open innovation, the new imperative for creating and profiting from technology*. Cambridge, MA: Harvard Business School.
- Christensen, C. (1997). *The innovator's dilemma: When new technologies cause great firms to fail*. Harvard Business Review Press.
- Cianci, P. J. (2014). *High definition television: The creation, development and implementation of HDTV technology*. McFarland.
- Cohen, W.M., & Levinthal, D. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, (35), 128-152.
- Cohoon, J., Hedge, S., Martin, W., & Richards, D. (1987). Punctuated equilibria: A parallel genetic algorithm. *Proceedings of the 2nd international conference on genetic algorithms* (pp. 28-31). Hillsdale, NJ: Lawrence Erlbaum Associates, Publishers.

References

- Coleman, J. S. (1988). Social capital in the creation of human capital. *American Journal of Sociology*, (94), S95-S120.
- Colovic, A., & Cartier, M. (2007). Appropriability, proximity routines and innovations. Paper presented at the *Druid Conference*.
- Cowan, R., Jonard, N., & Zimmermann, J. B. (2006). Evolving networks of inventors. *Journal of Evolutionary Economics*, (16, 1-2), 155-174.
- 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(2), 320-342.
- Cowan, R., Jonard, N., & Özman, M. (2004). Knowledge dynamics in a network industry. *Technological Forecasting and Social Change*, 71(5), 469-484.
- Cowan, R., Jonard, N., & Zimmermann, J. (2007). Bilateral collaboration and the emergence of innovation networks. *Management Science*, (53-7), 1051-1067.
- Czaszar, F. A., & Siggelkow, N. (2010). How much to copy? determinants of effective imitation breadth. *Organization Science*, 21(3), 661-676.
- Cusumano, M. A., & Gawer, A. (2002). The elements of platform leadership. *MIT Sloan Management Review*, 43(3), 51-58.

D

- Darwin, C. (1859). *On the origin of species by means of natural selection*. London: Murray.
- Dasgupta, P., & David, P. A. (1994). Toward a new economics of science. *Research Policy*, 23(5), 487-521.
- Desai, D. A. (2010). Co-creating learning: Insights from complexity theory. *The Learning Organization*, 17(5), 388-403.
- Dicken, P. (2015). *Global shift: Mapping the changing contours of the world economy* (7th ed.) SAGE Publications Ltd.
- Dosi, G. (1988). Sources, procedures, and microeconomic effects of innovation. *Journal of Economic Literature*, 26, 1120-1171.

E

- Elizalde, F. (2018). *Hype cycle for the connected home*. Stamford: Gartner. Retrieved from <https://www.gartner.com/document/3884568>
- Elmuti, D. (2003). The perceived impact of outsourcing on organizational performance. *American Journal of Business*, 18(2), 33-42.
- Ernst, D., & Kim, L. (2002). Global production networks, knowledge diffusion, and local capability formation. *Research Policy*, 31, 1417-1429.
- Ethiraj, S. K., & Posen, H. E. (2014). Do product architectures affect innovation productivity in complex product ecosystems? *Collaboration and competition in business ecosystems* (pp. 127-166). Emerald Group Publishing Limited.

F

- Farrell, J., & Katz, M. (2000). Innovation, rent extraction, and integration in systems markets. *The Journal of Industrial Economics*, 48(4), 413-432.
- 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(6), 465-487.
- Fleming, L., & Sorenson, O. (2001a). The dangers of modularity. *Harvard Business Review*, 79(8), 20-21.
- Fleming, L., & Sorenson, O. (2001b). Technology as a complex adaptive system: Evidence from patent data. *Research Policy*, 30(7), 1019-1039.
- Fleming, L., King III, C., & Juda, A. I. (2007a). Small worlds and regional innovation. *Organization Science*, 18(6), 938-954.
- Fleming, L., Mingo, S., & Chen, D. (2007b). Collaborative brokerage, generative creativity, and creative success. *Administrative Science Quarterly*, 52(3), 443-475.

- Franzoni, C., & Sauermann, H. (2014). Crowd science: The organization of scientific research in open collaborative projects. *Research Policy*, 43(1), 1-20.
- Freeman, C. (1987). Technical innovation, diffusion, and long cycles of economic development. *The long-wave debate* (pp. 295-309). Berlin: Springer.
- Freeman, L. (1977). A set of measures of centrality based upon betweenness. *Sociometry* 40, 35-41.
- Frenken, K. (2006a). A fitness landscape approach to technological complexity, modularity, and vertical disintegration. *Structural Change and Economic Dynamics*, 17(3), 288-305.
- Frenken, K. (2006b). *Innovation, evolution and complexity theory*. Cheltenham; UK: Edward Elgar Publishing.
- Frenken, K., Izquierdo R. L., & Zeppini, P. (2012). Branching innovation, recombinant innovation, and endogenous technological transitions. *Environmental Innovation and Societal Transitions*, 4, 25-35.
- Frenken, K., & Mendritski, S. (2012). Optimal modularity: A demonstration of the evolutionary advantage of modular architectures. *Journal of Evolutionary Economics*, 22(5), 935-956.
- Frenken, K., Van Oort, F., & Verburg, T. (2007). Related variety, unrelated variety and regional economic growth. *Regional Studies*, 41(5), 685-697.
- Funk, R.J. (2014). Making the most of where you are: Geography, networks, and innovation in organizations. *Academy of Management Journal*, 57(1), 193-222.

G

- Gatti, C., Volpe, L., & Vagnani, G. (2015). Interdependence among productive activities: Implications for exploration and exploitation. *Journal of Business Research*, 68(3), 711-722.
- Gavetti, G., & Levinthal, D. (2000). Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly*, 45(1), 113-137.
- Gawer, A. (2014). Bridging differing perspectives on technological platforms: Toward an integrative framework. *Research Policy*, 43(7), 1239-1249.
- Gawer, A., & Cusumano, M. A. (2008). How companies become platform leaders. *MIT Sloan Management Review*, 49(2), 28-35.
- Gawer, A., & Cusumano, M. A. (2014a). Industry platforms and ecosystem innovation. *Journal of Product Innovation Management*, 31(3), 417-433.
- Gawer, A., & Cusumano, M. A. (2014b). Platforms and innovation. In M. Dodgson, D. Gann, N. Phillips (Eds.), *The oxford handbook of innovation management*. Oxford: Oxford University Press.
- Gawer, A., & Henderson, R. (2007). Platform owner entry and innovation in complementary markets: Evidence from intel. *Journal of Economics & Management Strategy*, 16(1), 1-34.
- Geels, F. W. (2002). Technological transitions as evolutionary reconfiguration processes: A multi-level perspective and a case-study. *Research Policy* 31(8-9), 1257-1274.
- Ghemawat, P., & Levinthal, D. (2008). Choice interactions and business strategy. *Management Science*, 54(9), 1638-1651.
- Gilsing, V., Nootboom, B., Vanhaverbeke, 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(10), 1717-1731.
- Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Reading, MA: Addison-Wesley.
- Goldberg, D. E., Kargupta, H., Horn, J., & Cantu-Paz, E. (1995). *Critical deme size for serial and parallel genetic algorithms*. Urbana-Champaign: University of Illinois.
- Grabher, G., & Stark, D. (1997). Organizing diversity: Evolutionary theory, network analysis and postsocialism. *Regional Studies* 31(5), 533-544.
- Ganco, M., & Hoetker, G. (2009). NK modeling methodology in the strategy literature: Bounded search on a rugged landscape. *Research methodology in strategy and management* (pp. 237-268) Emerald Group Publishing Limited.
- Granovetter, M. (1985). Economic action and social structure: The problem of embeddedness. *American Journal of Sociology*, 91(3), 481-510.
- Gubbins, C., & Dooley, L. (2014). Exploring social network dynamics driving knowledge management for innovation. *Journal of Management Inquiry*, 23, 162-185.
- Gupta, A. K., Smith, K. G., & Shalley, C. E. (2006). The interplay between exploration and exploitation. *Academy of Management Journal*, 49, 693-706.

H

- Haller, J. B. A., Bullinger, A. C., & Möslin, K. M. (2011). Innovation contests: An IT-based tool for innovation management. *Business & Information Systems Engineering*, 3, 103-106.
- Hamming, R. W. (1950). Error detecting and error correcting codes. *Bell System Technical Journal*, 29(2), 147-160.
- Hansen, M. T. (1999). The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. *Administrative Science Quarterly*, 44, 82-111.
- Hart, O., & Moore, J. (1990). Property rights and the nature of the firm. *Journal of Political Economy*, 98(6), 1119-1158.
- Henderson, R. M., & Clark, K. B. (1990). Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. *Administrative Science Quarterly*, 35(1), 9-30.
- Hite, J., & Hesterly, W. S. (2001). The evolution of firm networks: From emergence to early growth of the firm. *Strategic Management Journal*, (22-3), 275-286.
- Holland, J. H. (1975). *Adaptation in natural and artificial systems*. Ann Arbor, MI: The University of Michigan Press.
- Holland, J. H. (1992). *Adaptation in natural and artificial systems*. Cambridge, MA: The MIT Press.
- Huber, F. (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(9), 1169-1182.

I

- Iansiti, M., & Levien, R. (2004). Strategy as ecology. *Harvard Business Review*, 82(3), 68-81.

J

- Jacobides, M. G., Cennamo, C., & Gawer, A. (2018). Towards a theory of ecosystems. *Strategic Management Journal*, 39(8), 2255-2276.
- Jain, S., & Krishna, S. (2001). A model for the emergence of cooperation, interdependence, and structure in evolving networks. *PNAS*, 98(2), 543-547.
- Janssen, M. J., & Frenken, K. (2019). Cross-specialisation policy: Rationales and options for linking unrelated industries. *Cambridge Journal of Regions, Economy and Society*, 12(2), 195-212.
- Joskow, P. L. (2008). Lessons learned from the electricity market liberalization. *The Energy Journal*, (Special Issue: The Future of Electricity: Papers in Honor of David Newbery), 9-42.

K

- Kapoor, R. (2018). Ecosystems: Broadening the locus of value creation. *Journal of Organization Design*, 7(1), 12.
- Kapoor, R., & Agarwal, S. (2017). Sustaining superior performance in business ecosystems: Evidence from application software developers in the iOS and android smartphone ecosystems. *Organization Science*, 28(3), 531-551.
- Kauffman, S.A. (1995). *At home in the universe*. Oxford: Oxford University Press.
- Kauffman, S. A. (1993). *The origins of order. self-organization and selection in evolution*. New York and Oxford: Oxford University Press.
- Kauffman, S. A., & Johnsen, S. (1991). Coevolution to the edge of chaos: Coupled fitness landscapes, poised states, and coevolutionary avalanches. *Journal of Theoretical Biology*, 149(4), 467-505.
- Kauffman, S. A., Lobo, J., & Macready, W. G. (2000). Optimal search on a technology landscape. *Journal of Economic Behavior & Organization*, 43(2), 141-166.
- Kauffman, S., & Macready, W. (1995). Technological evolution and adaptive organizations: Ideas from biology may find applications in economics. *Complexity*, 1(2), 26-43.
- Kenney, M., & Zysman, J. (2016). The rise of the railway economy. *Issues in Science and Technology*, 32(3), 61.
- Kleinová, E. (2016). Does liberalization of the railway industry lead to higher technical effectiveness? *Journal of Rail Transport Planning & Management*, 6(1), 67-76.
- Klepper, S. (1997). Industry life cycles. *Industrial and Corporate Change*, 6(1), 145-182.
- Kotabe, M., & Mol, M. J. (2009). Outsourcing and financial performance: A negative curvilinear effect. *Journal of Purchasing and Supply Management*, 15(4), 205-213.

L

- Lane, P.J., & Lubatkin, M. (1998). Relative absorptive capacity and interorganizational learning. *Strategic Management Journal*, 19(5), 461-477.
- 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(3), 871-897.
- Lavie, D., Stettner, U., & Tushman, M. L. (2010). Exploration and exploitation within and across organizations. *The Academy of Management Annals*, 4(1), 109-155.
- Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, 52(4), 667-694.
- Lazonick, W., Mazzucato, M., & Tulum, Ö. (2013). Apple's changing business model: What should the world's richest company do with all those profits? Paper presented at the *Accounting Forum*, 37(4), 249-267.
- Lee, E., Lee, J., & Lee, J. (2006). Reconsideration of the winner-take-all hypothesis: Complex networks and local bias. *Management Science*, 52(12), 1838-1848.
- LeHong, H., Burton, B., Bloesch, M., Lowendahl, J., & Moyer, K. (2019). *Eight ways ecosystems supercharge business models*. Stamford: Gartner. Retrieved from <https://www.gartner.com/document/3846964>
- Levinthal, D. A. (1997). Adaptation on rugged landscapes. *Management Science*, 31(7), 934-950.
- Levinthal, D. A., & Workiewicz, M. (2018). When two bosses are better than one: Nearly decomposable systems and organizational adaptation. *Organization Science*, 29(2), 207-224.
- Leviton, B., Lobo, J., Schuler, R., & Kauffman, S. A. (2002). Evolution of organizational performance and stability in a stochastic environment. *Computational & Mathematical Organization Theory*, 8(4), 281-313.
- Lieberman, M. B., & Asaba, S. (2006). Why do firms imitate each other? *Academy of Management Review*, 31(2), 366-385.
- Lissoni, F. (2001). Knowledge codification and the geography of innovation: The case of Brescia mechanical cluster. *Research Policy*, 30(9), 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(3), 195-220.
- Luo, J. (2018). Architecture and evolvability of innovation ecosystems. *Technological Forecasting and Social Change*, 136, 132-144.

M

- Mäkinen, S. J., & Dedehayir, O. (2014). Business ecosystems' Evolution—An ecosystem clockspeed perspective. *Collaboration and competition in business ecosystems* (pp. 99-125) Emerald Group Publishing Limited.
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2 (1), 71-87.
- Marion, R. (1999). *The edge of organization: Chaos and complexity theories of formal social systems*. Thousand Oaks, CA: Sage.
- Markard, J., & Truffer, B. (2006). Innovation processes in large technical systems: Market liberalization as a driver for radical change? *Research Policy*, 35(5), 609-625.
- Marsili, O. (2002). Technological regimes and sources of entrepreneurship. *Small Business Economics*, 19(3), 217-231.
- Martinsuo, M. (2013). Project portfolio management in practice and in context. *International Journal of Project Management*, 31(6), 794-803.
- McIntyre, D. P., & Subramaniam, M. (2009). Strategy in network industries: A review and research agenda. *Journal of Management*, 35(6), 1494-1517.
- McKelvey, B. (1999). Avoiding complexity catastrophe in coevolutionary pockets: Strategies for rugged landscapes. *Organization Science*, 10(3), 294-321.
- McNerney, J., Doyne Farmer, J., Redner, S., & Trancik E., J. (2011). Role of design complexity in technology improvement. *PNAS*, 108(22), 9008-9013.
- Milgrom, P., & Roberts, J. (1990). Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica: Journal of the Econometric Society*, 1255-1277.
- Miller, K. D., Zhao, M., & Calantone, R. J. (2006). Adding interpersonal learning and tacit knowledge to march's exploration-exploitation model. *Academy of Management Journal*, 49(4), 709-722.
- Miller, R., Hobday, M., Leroux-Demers, T., & Olleros, X. (1995). Innovation in complex systems industries: The case of flight simulation. *Industrial and Corporate Change*, 4(2), 363-400.

References

- Mitton, C., Adair E., C., McKenzie, E., Patten B., S., & Wayne Perry, B. (2007). Knowledge transfer and exchange: Review and synthesis of the literature. *The Millbank Quarterly*, 85(4), 729-768.
- Moore, J. F. (1993). Predators and prey: A new ecology of competition. *Harvard Business Review*, 71(5), 75-86.
- Moran, P., Simoni, M., & Vagnani, G. (2011). Becoming the best: By beating or ignoring the best? toward an expanded view of the role of managerial selection in complex and turbulent environments. *Journal of Management & Governance*, 15(3), 447-481.
- Mowery, D. C., Oxley, J. E., & Silverman, B. S. (1998). Technological overlap and interfirm cooperation: Implications for the resource-based view of the firm. *Research Policy*, (27), 507-523.
- Murmann, J. P., & Frenken, K. (2006). Toward a systematic framework for research on dominant designs, technological innovations, and industrial change. *Research Policy*, 35(7), 925-952.
- Muthukrishna, M., & Henrich, J. (2016). Innovation in the collective brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1690), 20150192.

N

- Nash, J. (1951). Non-cooperative games. *Annals of Mathematics*, 286-295.
- Nelson, R., & Winter, S. (1982). *An evolutionary theory of economic change*. Cambridge, MA: Harvard University Press.
- Newman, M. E. J. (1999). Small worlds - the structure of social networks. *Santa Fe Institute Working Papers*, (99-12-080), <http://www.santafe.edu/media/workingpapers>.
- Newman, M. E. (2003). The structure and function of complex networks. *SIAM Review*, 45(2), 167-256.
- Nonaka, I. (1991). The knowledge-creating company. *Harvard Business Review*, July-August, 96-104.
- Nooteboom, B. (2000). Learning by interaction: Absorptive capacity, cognitive distance and governance. *Journal of Management and Governance*, 4(1-2), 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(7), 1016-1034.
- Norman, D. A., & Verganti, R. (2014). Incremental and radical innovation: Design research vs. technology and meaning change. *Design Issues*, 30(1), 78-96.

O

- Ordanini, A., Rubera, G., & DeFillippi, R. (2008). The many moods of inter-organizational imitation: A critical review. *International Journal of Management Reviews*, 10(4), 375-398.
- Oxley, J. E., & Sampson, R. C. (2004). The scope and governance of international R&D alliances. *Strategic Management Journal*, 25(8-9), 723-749.
- Ozman, M. (2009). Inter-firm networks and innovation: A survey of literature. *Economic of Innovation and New Technology*, 18(1), 39-67.

P

- Peltoniemi, M., & Vuori, E. (2004). Business ecosystem as the new approach to complex adaptive business environments. Paper presented at the *Proceedings of eBusiness Research Forum*, 2, 267-281.
- Pettey, C. B., Leuze, M. R., & Grefenstette, J. J. (1987). Parallel genetic algorithm. Paper presented at the *Genetic Algorithms and their Applications: Proceedings of the Second International Conference on Genetic Algorithms: July 28-31, 1987 at the Massachusetts Institute of Technology, Cambridge, MA*.
- Phelps, C., Heidl, R., & Wadhwa, A. (2012). Knowledge, networks, and knowledge networks: A review and research agenda. *Journal of Management*, 38(4), 1115-1166.
- Porter, M., & Siggelkow, N. (2008). Contextuality within activity systems and sustainability of competitive advantage. *Academy of Management Perspectives*, 22(2), 34-56.
- Posen, H. E., Lee, J., & Yi, S. (2013). The power of imperfect imitation. *Strategic Management Journal*, 34(2), 149-164.
- Posen, H. E., & Martignoni, D. (2018). Revisiting the imitation assumption: Why imitation may increase, rather than decrease, performance heterogeneity. *Strategic Management Journal*, 39(5), 1350-1369.
- Postmes, T., Spears, R., & Cihangir, S. (2001). Quality of decision making and group norms. *Journal of Personality and Social Psychology*, 80(6), 918-930.
- 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(1), 116-145.

- Press, K. (2008). Divide to conquer? limits to the adaptability of disintegrated, flexible specialization clusters. *Journal of Economic Geography*, 8(4), 565-580.
- Puranam, P., Stieglitz, N., Osman, M., & Pillutla, M. (2015). Modelling bounded rationality in organizations: Progress and prospects. *Academy of Management Annals*, 9(1), 337-392.
- Pyka, A. (2002). Innovation networks in economics: From the incentive-based to the knowledge-based approaches. *European Journal of Innovation Management*, 5(3), 152-163.

Q

- Querbes, A., & Frenken, K. (2012). Why first movers may fail: Global versus sequential improvement of complex technological artefacts. *Proceedings of the European conference on complex systems 2012* (pp. 751) Springer International Publishing.
- Querbes, A., & Frenken, K. (2017). Evolving user needs and late-mover advantage. *Strategic Organization*, 15(1), 67-90.

R

- Raisch, S., Birkinshaw, J., Probst, G., & Tushman, M. L. (2009). Organizational ambidexterity: Balancing exploitation and exploration for sustained performance. *Organization Science*, 20(4), 685-695.
- Rivkin, J. W. (2000). Imitation of complex strategies. *Management Science*, 46(6), 824-844.
- Rivkin, J. W., & Siggelkow, N. (2003). Balancing search and stability: Interdependencies among elements of organizational design. *Management Science*, 49(3), 290-311.
- Rochet, J. C., & Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4), 990-1029.
- Rogers, A., & Prügél-Bennett, A. (1999). Genetic drift in genetic algorithm selection schemes. *IEEE Transactions on Evolutionary Computation*, 3(4), 298-303.
- Roundy, P. T., Brockman, B. K., & Bradshaw, M. (2017). The resilience of entrepreneurial ecosystems. *Journal of Business Venturing Insights*, 8, 99-104.
- Rowley, T., Behrens, D., & Krackhardt, D. (2000). Redundant governance structures: An analysis of structural and relational embeddedness in the steel and semiconductor industries. *Strategic Management Journal*, 21(3), 369-386.
- Russell, M. G., & Smorodinskaya, N. V. (2018). Leveraging complexity for ecosystemic innovation. *Technological Forecasting and Social Change*, 136, 114-131.

S

- Sabidussi, G. (1966). The centrality index of a graph. *Psychometrika*, 31, 581-603.
- 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(1), 54-75.
- Sawhney, M., Verona, G., & Prandelli, E. (2005). Collaborating to create: The internet as a platform for customer engagement in product innovation. *Journal of Interactive Marketing*, 19(4), 4-17.
- Schepers, J., Schnell, R., & Vroom, P. (1999). How Siemens bridges the innovation gap. *Research-Technology Management*, 42, 26-31.
- Schilling, M., & Phelps, C. (2007). Interfirm collaboration networks: The impact of small world connectivity on firm innovation. *Management Science*, 53(7), 1113-1126.

References

- Schumpeter, J. A. (1939). *Business cycles: A theoretical, historical and statistical analysis of the capitalist process*. New York: McGraw-Hill.
- Seetharaman, G., Lakhota, A., & Blasch, P. E. (2006). Unmanned vehicles come of age: The DARPA grand challenge. *Computer*, 39(12)
- Shah, S. K. (2006). Motivation, governance, and the viability of hybrid forms in open source software development. *Management Science*, 52(7), 1000-1014.
- Silverberg, G., & Verspagen, B. (2005). A percolation model of innovation in complex technology spaces. *Journal of Economic Dynamics and Control*, 29, 225-244.
- Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6), 467-482.
- Simon, H. A. (1996). *The sciences of the artificial* (3rd ed.). Cambridge: The MIT Press.
- Solis-Molina, M., Hernández-Espallardo, M., & Rodríguez-Orejuela, A. (2018). Performance implications of organizational ambidexterity versus specialization in exploitation or exploration: The role of absorptive capacity. *Journal of Business Research*, 91, 181-194.
- Sorenson, O., Rivkin, J. W., & Fleming, L. (2006). Complexity, networks and knowledge flow. *Research Policy*, 35(7), 994-1017.
- Stam, C. J. (2004). Functional connectivity patterns of human magnetoencephalographic recordings: A 'small-world' network? *Neuroscience Letters*, 355, 25-28.
- 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(3), 1250-1259.
- Stettner, U., & Lavie, D. (2014). Ambidexterity under scrutiny: Exploration and exploitation via internal organization, alliances, and acquisitions. *Strategic Management Journal*, 35(13), 1903-1929.
- Stewart, M. (2001). *The coevolving organization*. Rutland: Decomplexity Associates LTD.

T

- Tansley, A. G. (1935). The use and abuse of vegetational concepts and terms. *Ecology*, 16, 284-307.
- Tedeschi, G., Mazloumian, A., Gallegati, M., & Helbing, D. (2012). Bankruptcy cascades in interbank markets. *PLoS One*, 7(12), e52749.
- Tee, R., & Gawer, A. (2009). Industry architecture as a determinant of successful platform strategies: A case study of the i-mode mobile internet service. *European Management Review*, 6(4), 217-232.
- Teece, D. J. (2018). Business models and dynamic capabilities. *Long Range Planning*, 51(1), 40-49.
- Ter Wal, Anne L. J. (2013). The dynamics of the inventor network in German biotechnology: Geographic proximity versus triadic closure. *Journal of Economic Geography*, 14(3), 1-32.
- Terwiesch, C., & Xu, Y. (2006). Innovation contests, open innovation, and multiagent problem solving. *Management Science*, 54(9), 1529-1543.
- Thierens, D., & Goldberg, D. (1994). Elitist recombination: An integrated selection recombination GA. Paper presented at the *Proceedings of the First IEEE Conference on Evolutionary Computation. IEEE World Congress on Computational Intelligence*, 508-512.
- Tsai, C. (2018). The timing of fostering complementary innovation: Exploring the antecedent of industry platform emergence. *Technology Analysis and Strategic Management*, 30(10), 1121-1135.
- Tushman, M. L., & Anderson, P. (1986). Technological discontinuities and organizational environments. *Administrative Science Quarterly*, 439-465.
- Tushman, M., & Nadler, D. (1986). Organizing for innovation. *California Management Review*, 28(3), 74-92.

U

- Uotila, J. (2017). Punctuated equilibrium or ambidexterity: Dynamics of incremental and radical organizational change over time. *Industrial and Corporate Change*, 27(1), 131-148.
- Uzzi, B. (1996). The sources and consequences of embeddedness for the economic performance of organizations: The network effect. *American Sociological Review*, 61(4), 674-698.
- Uzzi, B., Amaral, L., & Reed-Tsochas, F. (2007). Small-world networks and management science research: A review. *European Management Review*, 4, 77-91.

V

- 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(5), 1094-1107.
- Van Valen, L. (1977). The red queen. *The American Naturalist*, 111(980), 809-810.
- Verspagen, B., & Duysters, G. (2004). The small worlds of strategic technology alliances. *Technovation*, 24(7), 563-571.
- Vidgen, R., & Bull, L. (2011). Applications of Kauffman's coevolutionary NKCS model to management and organization studies. *Complexity and management* (pp. 201-219). Thousand Oaks, CA: Sage.
- von Hippel, E. (2005). *Democratizing innovation*. Cambridge, MA: MIT Press.

W

- Wareham, J., Fox, P. B., & Giner, J. L. C. (2014). Technology ecosystem governance. *Organization Science*, 25(4), 1195-1215.
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393, 440-442.
- Weinberger, A., Stegmann, K., & Fischer, F. (2007). Knowledge convergence in collaborative learning: Concepts and assessment. *Learning and Instruction*, 17(4), 416-426.
- Welter, C., & Kim, S. (2018). Effectuation under risk and uncertainty: A simulation model. *Journal of Business Venturing*, 33(1), 100-116.
- West, G. (2017). *Scale: The universal laws of growth, innovation, sustainability, and the pace of life in organisms, cities, economies, and companies*. New York: Penguin Press.
- West, J., & Gallagher, S. (2006). Patterns of open innovation in open source software. In H. Chesbrough, W. Vanhaverbeke & J. West (Eds.), *Open innovation: Researching a new paradigm*. Oxford University Press Oxford.
- Wilden, R., Hohberger, J., Devinnney, T. M., & Lavie, D. (2018). Revisiting James March (1991): Whither exploration and exploitation? *Strategic Organization*, 16(3), 352-369.
- Wilensky, U. (1999). Netlogo. center for connected learning and computer-based modeling. Retrieved from <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., & Nootboom, B. (2005). Empirical tests of optimal cognitive distance. *Journal of Economic Behavior & Organization*, 58(2), 277-302.

Y

- Yoffie, D. B., & Kwak, M. (2006). With friends like these: The art of managing complementors. *Harvard Business Review*, 84(9), 88-98.

Z

- Zeppini, P., Frenken, K., & Kupers, R. (2014). Thresholds models of technological transitions. *Environmental Innovation and Societal Transitions*, 11, 54-70.

Summary

Human and societal progress depend on our capability to innovate. Over the centuries, many problems and challenges in welfare, health, climate, sustainability and other areas, have been addressed by countless innovations. Over and over again, individuals and organizations have applied state-of-the-art knowledge and technologies to find novel solutions.

Increasingly, however, innovation is no longer done in isolation. Ever more globalization, digitization and complexity both enable and require a more collective approach to innovation, taking place in what is known as a 'business ecosystem'. All organizations are part of one or more of such ecosystems, in which firms and other actors collaborate, compete or interact in a myriad of other ways (Adner, 2012).

Considering, analyzing, understanding and improving innovation from an ecosystem perspective is therefore of critical importance. In the work that is presented in this thesis, this ecosystem perspective, complemented with concepts and insights from complexity science and evolution science, is taken as a general approach to extend and enrich extant work in innovation studies and management science. As much of the literature in this area is focused on empirical work, an important goal of the current work has been to underpin this by a theoretical and analytical research approach, using modeling and simulations. More in particular, in the current work a distinction is made between a number of more specific ecosystem perspectives: contest innovation, market innovation and co-innovation.

In contest innovations, described in chapter 2, firms organize an innovation competition between their internal units or external contenders. Instead of doing the innovation internally, using a common, single approach, the key idea is that a contest provides a way to address the same innovation challenge by multiple approaches in parallel, increasing the likelihood of finding a (better) solution. However, this does not imply that contestants should not share knowledge, as this will help all participants to avoid dead-ends and to identify promising avenues for exploration. At the same time, overly knowledge sharing should be avoided too, to prevent that all contenders converge on the same approach, perhaps prematurely and destroying the diversity that was sought in the first place. Simulating the process of contest innovation, including several variants and frequencies of knowledge sharing, revealed that finding a balance between too little and too much knowledge sharing is indeed critical. In other words, the right balance yields best performance, even more so with growing difficulty of the challenge at hand.

In market innovation, firms do their own innovation, but may also try to imitate innovations made by other, leading firms. However, mistakes are easily made during imitation, especially in case of products of high complexity. A key finding, as presented in chapter 3, holds that successful imitation depends on both social proximity and cognitive proximity between firms, with social proximity between employees of different firms being an important substitute for cognitive distance or a lack of shared knowledge that limits absorptive capacity. Nevertheless, too much social or cognitive distance will render imitation invariable, especially when complexity is higher. On the other hand, too little distance will render imitation useless, as proximate firms will tend to become more similar. In addition, it was found that the structure of the social network between firms is

another relevant factor in this context. Again a balance must be found here, between on the one hand too much clustering of firms, providing insufficient new innovation ideas, and on the other hand too much randomness and short network distances, leading to premature convergence. Small-world networks provide such a balance, offering an explanation as to why such networks are so often found in empirical work.

In co-innovation, firms collaborate in their ecosystem to do innovation, with each firm focusing on its own components. Together, components of different firms can be combined into a value creating product or service. This implies that a firm not only has to manage the dependencies between its own components – internal complexity – but also has to manage the dependencies and find complementarities between its own components and those of other firms – external complexity. Chapter 4 describes a model that captures both internal and external complexity, leading to the emergence of different ecosystem regimes. These regimes are either rigid or chaotic, with the so called ‘edge of chaos’ regime in between. Rigid regimes are characterized by high levels of internal complexity, making it very hard for firms to innovate, as each change requires a resynchronization with many other components. In contrast, a chaotic regime is characterized by high levels of external complexity, causing firms to continuously respond to changes in other firms, rendering innovations quickly useless. It is only in or near the edge of chaos regime that firms manage to balance stability for sustainable innovation with adaptability to respond to external changes. However, an important finding in the current work holds that the occurrence of these regimes is also dependent on the structure of dependencies between firms in their ecosystem. In general, clustering will push an ecosystem more towards a rigid regime, while a more random network structure will push towards a chaotic regime. Again, small-world networks seem to offer the best of both worlds. The latter was also found to be the case when considering the resilience of an ecosystem in responding to shocks, such as the disappearance of an existing firm and the introduction of a new firm. Such shocks easily turn into performance degrading avalanches in a chaotic regime, while overly rigid regimes will fail to use shocks as an opportunity to escape from local optima (and Nash equilibria) to further innovation.

From a more integral perspective, and building on the results of the previous chapters, chapter 5 introduces the ecosystem complexity model. Using the notions of internal and external complexity as dimensions, this model identifies four main ecosystem roles. These roles are *completer*, firms that are largely independent in their innovation, *composers*, firms that offer a platform or infrastructure that can be used by *complementors*, that complement the platform and infrastructure with their own components, thus creating useful products for their clients. Finally, *connectors* are firms that focus on offering adapters or brokering services to enable the combination of products or services of other firms. Innovation strategy, then, is very much about choosing a firm’s position in the ecosystem in terms of one of the four roles, and employing strategies to move a firm from its current to a desired position. These strategies include (dis)intermediation, (de)modularization and insourcing/outsourcing, changing a firm’s internal and external dependencies. Another key result from the application of the ecosystem complexity model is its use to understand the lifecycle of industries. In summary, many new industries are populated by completers, by lack of other firms. Connectors find a role to combine products of different

completers, or of composers that start to offer platforms, applying emerging standards, thus enabling the rise of complementors.

In terms of theoretical implications, the current work has further demonstrated the value of considering innovation from an ecosystem perspective. Simulations based on evolution theory of models based on complexity theory, have added to theoretical insights to underpin empirical work and to explain multiple stylized facts. To accomplish this, seminal work on the NK-model (Kauffman, 1993), the NKCS-model (Kauffman & Johnsen, 1991), network structures (Watts & Strogatz, 1998) and the product lifecycle (Abernathy and Utterback, 1978) has been combined and extended.

More practical implications for innovation managers and policy makers are manifold. These include the need to include the ecosystem perspective in strategic decision making. Firms should build and maintain awareness about their ecosystem, its network structure, their own position and the positions of other firms. Policy makers and firms, as far as possible, should use these notions also to influence their ecosystem structure and positions. For example by actively influencing knowledge sharing and social networks, by balancing clustering versus network short-cuts or by applying strategies such as outsourcing, intermediation or modularization, or their reverse, depending on their ecosystem-aware innovation strategy.



Samenvatting

Menselijke en maatschappelijke vooruitgang zijn afhankelijk van ons vermogen te innoveren. Gedurende de eeuwen zijn vele problemen en uitdagingen op het vlak van welzijn, gezondheid, klimaat, duurzaamheid en andere gebieden aangepakt met talloze innovaties. Telkens weer zijn individuen en organisaties op basis van de laatste kennis en techniek op zoek gegaan naar nieuwe oplossingen.

Echter, in toenemende mate vindt innovatie niet meer plaats in isolatie. De almaar verdergaande globalisering, digitalisering en complexiteit maken een meer collectieve benadering van innovatie zowel mogelijk als noodzakelijk, in de context van zogenaamde 'business ecosystemen'. Elke organisatie is onderdeel van één of meer van dergelijke ecosystemen, waarin bedrijven en andere actoren samenwerken, concurreren of op talloze andere manieren met elkaar interacteren (Adner, 2012).

Het beschouwen, analyseren, begrijpen en verbeteren van innovatie vanuit het perspectief van ecosystemen is daarmee van groot belang. In het in deze thesis gepresenteerde werk is dit ecosysteem perspectief, aangevuld met concepten en inzichten uit de complexiteitswetenschap en evolutieleer, de gemeenschappelijke benadering om bestaand werk in innovatiestudies en managementwetenschap uit te breiden en te verrijken. Aangezien veel van de bestaande literatuur in dit gebied geconcentreerd is op empirisch werk, is een belangrijk doel van het huidige werk om dit te onderbouwen met een meer theoretische en analytische benadering, gebruikmakend van modellen en simulaties. Meer in het bijzonder is in het huidige werk een onderscheid gemaakt naar een aantal meer specifieke ecosysteem perspectieven: competitie innovatie, markt innovatie en co-innovatie.

In een innovatiewedstrijd, zoals beschreven in hoofdstuk 2, organiseren bedrijven een competitie tussen hun interne onderdelen of externe deelnemers. In plaats van het intern uitvoeren van innovaties met een enkele, gezamenlijke aanpak, is het kernidee dat in een competitie dezelfde innovatie uitdaging vanuit verschillende invalshoeken parallel wordt aangegaan, daarmee de kans vergroten op het vinden van een (betere) oplossing. Dit betekent echter niet dat deelnemers geen kennis zouden moeten delen, aangezien dit alle deelnemers kan helpen om eventuele doodlopende straten te vermijden en andere meer veelbelovende routes te vinden. Tegelijkertijd is een teveel aan kennisdeling ongewenst, aangezien de deelnemers dan wellicht voorbarig allen dezelfde aanpak gaan volgen, wat ten koste zou gaan van de juist gezochte diversiteit. Uit simulaties van competitie innovatie, inclusief verschillende varianten en frequenties van kennisdeling, is gebleken dat het vinden van een balans tussen te veel en te weinig kennisdeling inderdaad van kritiek belang is. Met andere woorden, het vinden van de juiste balans leidt tot de beste prestaties, vooral naarmate het innovatieprobleem ingewikkelder wordt.

In markt-innovatie doet elk bedrijf aan eigen innovatie. Daarnaast kunnen bedrijven proberen om de innovaties van andere, vooroplopende bedrijven te imiteren. Fouten zijn echter snel gemaakt bij het imiteren, vooral in het geval van producten met een hoge complexiteit. Een belangrijke bevinding, gepresenteerd in hoofdstuk 3, is dat succesvolle imitatie afhankelijk is van zowel sociale als cognitieve nabijheid tussen bedrijven. Daarbij is de sociale nabijheid tussen medewerkers van verschillende bedrijven een belangrijk substituuat voor meer cognitieve afstand, ofwel

een gebrek aan gedeelde kennis die de absorptiecapaciteit belemmert. Niettemin zal te veel sociale of cognitieve afstand het imiteren onmogelijk maken, vooral bij hogere complexiteit. Aan de andere kant zal een te kleine afstand het imiteren juist nutteloos maken, aangezien nabije bedrijven al veel op elkaar lijken. In aanvulling hierop is gevonden dat de structuur van het sociale netwerk tussen bedrijven een andere belangrijke factor is in deze context. Opnieuw is het van belang een balans te vinden, tussen aan de ene kant een teveel aan clustering van bedrijven, waaruit te weinig nieuwe innovatie ideeën voortkomen, en aan de andere kant een te willekeurig netwerk met korte afstanden, leidend tot premature convergentie. *Small-world* netwerken bieden een dergelijke balans, hetgeen een verklaring biedt voor het vaak aantreffen van dergelijke netwerken in empirisch werk.

In co-innovatie werken bedrijven samen om te innoveren, waarbij elk bedrijf zich focust op de eigen componenten. Gezamenlijk combineren bedrijven deze componenten in een compleet product of dienst. Dit impliceert dat een bedrijf niet alleen complementariteiten en afhankelijkheden tussen de eigen componenten – interne complexiteit – moet managen, maar ook die tussen de eigen componenten en die van andere bedrijven – externe complexiteit. Hoofdstuk 4 beschrijft een model dat zowel interne als externe complexiteit representeert, leidend tot de emergentie van verschillende ecosysteem regimes. Deze regimes zijn ofwel rigide ofwel chaotisch, met het *edge of chaos* regime er tussenin. Rigide regimes worden gekenmerkt door een hoge mate van interne complexiteit die het lastig maakt voor bedrijven om te innoveren, omdat elke verandering in een component een re-synchronisatie vergt met vele andere componenten. Chaotische regimes daarentegen worden gekenmerkt door een hoge mate van externe complexiteit, waardoor bedrijven continu moeten reageren op veranderingen in andere bedrijven en bestaande innovaties snel nutteloos worden. Alleen in of nabij het *edge of chaos* regime vinden bedrijven de balans tussen stabiliteit voor duurzame innovatie en adaptatievermogen om te reageren op externe veranderingen. Een belangrijke bevinding in het huidige werk is dat het optreden van deze regimes ook afhankelijk is van de structuur van afhankelijkheden tussen bedrijven in hun ecosysteem. In het algemeen zal clustering een ecosysteem meer in de richting van een rigide regime duwen, terwijl een meer willekeurige structuur in de richting van een chaotisch regime zal duwen. Opnieuw bieden *small-world* netwerken het beste van twee werelden. Dit laatste blijkt ook het geval te zijn als het gaat om de veerkracht van ecosystemen om schokken op te vangen, zoals het verdwijnen van een bestaand of het ontstaan van een nieuw bedrijf. Zulke schokken leiden in chaotische regimes snel tot lawines van veranderingen en sterke vermindering van prestatieniveaus, terwijl in te rigide regimes juist de kans wordt gemist om schokken te gebruiken als een kans om uit lokale optima (en Nash-evenwichten) te ontsnappen en innovatie te bevorderen.

Vanuit een meer integraal perspectief en voortbouwend op de resultaten uit eerdere hoofdstukken, introduceert hoofdstuk 5 het ecosysteem complexiteit model. Met interne en externe complexiteit als dimensies, identificeert dit model vier hoofdrollen in het ecosysteem. Deze rollen zijn *completers*, met bedrijven die grotendeels onafhankelijk zijn in hun innovaties, *composers*, met bedrijven die een platform of infrastructuur bieden aan *complementors*, met bedrijven die hierop hun eigen componenten kunnen aanbieden, zodat nuttige producten of diensten

ontstaan. Tot slot zijn er *connectors*, met bedrijven die zich richten op het bieden van adapters of bemiddelingsdiensten om de combinatie mogelijk te maken van producten en diensten van andere bedrijven. Innovatiestrategie bestaat hiermee uit het kiezen van een ecosysteem positie voor een bedrijf in termen van één van deze vier rollen, en het toepassen van strategieën om een bedrijf van de huidige naar de gewenste positie te bewegen. Deze strategieën bestaan onder meer uit (dis)intermediatie, (de)modulariseren en *insourcing* / *outsourcing*, die alle een uitwerking hebben op de interne en externe complexiteit. Een ander belangrijk resultaat uit de toepassing van het ecosysteem complexiteit model betreft het beter begrijpen van de levenscyclus van industriesectoren. Samengevat worden nieuwe sectoren initieel vooral bevolkt door *completers*, bij gebrek aan andere bedrijven om mee samen te werken. *Connectors* vinden dan een rol om de producten te combineren van verschillende *completers*, of van *composers* die gebruik maken van standaarden, waarmee de opkomst van *complementors* mogelijk wordt.

In termen van theoretische implicaties heeft het huidige werk de waarde verder aangetoond om innovatie vanuit een ecosysteem perspectief te beschouwen. Simulaties gebaseerd op evolutieel en modellen gebaseerd op complexiteitstheorie, hebben bijgedragen aan de theoretische inzichten om empirisch werk te onderbouwen en meerdere gestileerde feiten te verklaren. Om dit te bereiken is het fundamentele werk op het vlak van het NK-model (Kauffman, 1993), het NKCS-model (Kauffman & Johnsen, 1991), netwerkstructuren (Watts & Strogatz, 1998) en productlevenscycli (Abernathy & Utterback, 1978) gecombineerd en uitgebreid.

De meer praktische implicaties voor innovatie managers en beleidsmakers zijn meervoudig. Deze betreffen de noodzaak om het ecosysteem perspectief mee te nemen in strategische besluitvorming. Bedrijven doen er goed aan om bewust te worden, en te blijven, van het belang van hun ecosysteem, de netwerkstructuur daarvan, de eigen positie daarin en die van andere bedrijven. Bedrijven en beleidsmakers zouden deze noties ook moeten gebruiken om invloed uit te oefenen op hun positie en de structuur van het ecosysteem, voor zover mogelijk., bijvoorbeeld met het actief beïnvloeden van kennisdeling, het balanceren van clustering met verkorte netwerkroutes of het toepassen van strategieën als outsourcing, intermediatie en modulariseren, of het tegenovergestelde, afhankelijk van hun ecosysteem-bewuste innovatiestrategie.



Acknowledgements

The story behind this thesis is one of a long and winding road. A story that before long got tangled up with other life events. Family circumstances, migrating to Norway and back, raising a family, not to mention job changes and other enterprises, all have contributed to a sometimes challenging, but above all rewarding and invigorating era in my life, with never a dull moment.

Now, at the end of this road and looking back, there are two answers to the question about my motivation to start and pursue this work. The first is about the sincere fascination for complexity science and the conviction that many real-world problems require its application. The second is about the desire to work at the edge of chaos that can be found where the relative stability of scientific theory meets the dynamics of practical innovation. Building bridges between the world of science and the world of business and society is not only essential to both, it is also my idea of fun. In my case, such bridges are about connecting the world of data and artificial intelligence with the world of innovation science.

During my work on this thesis, my employers Capgemini, Alliander and Gartner have provided support and motivation. The CIO office of Alliander offered me time and funding to attend several conferences, including my visit to the Santa Fe Institute. In addition, both the department of Industrial Engineering & Innovation Sciences at the Eindhoven University of Technology and the Copernicus Institute at Utrecht University, were kind enough to host me and facilitate my work.

Special thanks to Floor Alkemade, at Eindhoven University of Technology, for her support and valuable comments and to my promotor Koen Frenken, at Utrecht University, for his guidance, knowledge, patience and inspiration.



Curriculum Vitae

After his Master of Science graduation at Utrecht University in one of the first cohorts in Cognitive Artificial Intelligence, specializing in neural networks and genetic algorithms at the national research institute for mathematics and computer science in the Netherlands, Pieter den Hamer joined knowledge center CIBIT. Following up on his work during his studies as independent data scientist, Pieter became responsible for research, consulting and education in knowledge technology, business intelligence and analytics. With the merger of CIBIT and the Software Engineering Research Centre (SERC), he became partner and managing director. After being acquired by DNV-GL, Pieter moved to Oslo, Norway, where he was research director for the 'innovative IT' program in DNV Research & Innovation. Returning to the Netherlands, he worked in several roles at Capgemini and energy company Alliander, before joining Gartner Research & Advisory in his current role as senior director artificial intelligence. Over the years he has published several managerial books, articles and papers in applied science. In addition, he has been an active member of multiple EU smart city & society initiatives, standardization bodies, research alliances and other organizations, most recently of the Centre for Complex Systems Studies.



List of publications (selection)

- Andrews, W., **den Hamer, P. J.**, Brethenoux, E., & Karamouzis, F. (2018). Artificial intelligence maturity model. Research note in *artificial intelligence series*. Stamford, CT: Gartner.
- den Hamer, P. J.** (2005). *De organisatie van business intelligence*. In book series 'essential IT competences for managers'. Den Haag: Academic Service.
- den Hamer, P. J.** (2009). *Business intelligence*. Den Haag: Sdu Uitgevers.
- den Hamer, P. J.** (2014), Energie = informatie. *Business Information Magazine*, 12, 31-35.
- den Hamer, P.**, & Frenken, K. (2020). A network-based model of exploration and exploitation. *Journal of Business Research*.
- den Hamer, P. J.**, & Rollings, M. (2018). Business ecosystems and continuous intelligence – how to use external data for competitive advantage. Research note in *applied analytics series*. Stamford, CT: Gartner.
- den Hamer, P. J.**, & Schulte, W. R. (2019). When to automate or augment decision making. Research note in *artificial intelligence series*. Stamford, CT: Gartner.
- den Hamer, P. J.**, Schulte, W. R., & Oestreich, T. (2018). Building your continuous intelligence capability for digital transformation. Research note in *applied analytics series*. Stamford, CT: Gartner.
- den Hamer, P.**, & Skramstad, T. (2011). Autonomic service-oriented architecture for resilient complex systems. Paper presented at the *2011 IEEE 30th Symposium on Reliable Distributed Systems Workshops*, 62-66.
- Mullen, A., Revang, M., & **den Hamer, P. J.** (2019). Design principles of human-in-the-loop systems for control, performance and transparency of AI. Research note in *artificial intelligence series*. Stamford, CT: Gartner.
- Schulte, W. R., & **den Hamer, P. J.** (2019). How to use machine learning, business rules and optimization in decision management. Research note in *applied analytics series*. Stamford, CT: Gartner.
- Titze, C., Tohamy, N., **den Hamer, P. J.**, & Payne, T. (2019). Innovation insight for continuous intelligence in supply chains. Research note in *supply chain management series*. Stamford, CT: Gartner.