




Knowledge Diffusion and Innovation

MODELLING COMPLEX ENTREPRENEURIAL BEHAVIOURS



Piergiuseppe Morone • Richard Taylor



Knowledge Diffusion and Innovation

To Mariarosaria, PM

Knowledge Diffusion and Innovation

Modelling Complex Entrepreneurial
Behaviours

Piergiuseppe Morone

University of Foggia, Italy

Richard Taylor

Stockholm Environment Institute, Oxford, UK

Edward Elgar

Cheltenham, UK • Northampton, MA, USA

© Piergiuseppe Morone and Richard Taylor 2010

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical or photocopying, recording, or otherwise without the prior permission of the publisher.

Published by
Edward Elgar Publishing Limited
The Lypiatts
15 Lansdown Road
Cheltenham
Glos GL50 2JA
UK

Edward Elgar Publishing, Inc.
William Pratt House
9 Dewey Court
Northampton
Massachusetts 01060
USA

A catalogue record for this book
is available from the British Library

Library of Congress Control Number: 2009937898



Mixed Sources

Product group from well-managed
forests and other controlled sources
www.fsc.org Cert no. SA-COC-1565
© 1996 Forest Stewardship Council

ISBN 978 1 84720 916 0

Printed and bound by MPG Books Group, UK

Contents

<i>List of figures</i>	vii
<i>List of tables</i>	ix
PART I THEORY	
1 Introduction	3
2 Knowledge economy: old and new issues	7
3 Modelling knowledge and its diffusion patterns: a pathway towards complexity	24
4 Knowledge diffusion and innovation: an agent-based approach	47
PART II EMPIRICAL STUDIES AND MODEL VALIDATIONS	
5 Empirical studies on knowledge flows	85
6 Theoretical and applied methodologies of agent-based models	104
7 Validating the model of knowledge diffusion and innovation	125
8 Final remarks and future research	151
<i>References</i>	160
<i>Index</i>	173

Figures

2.1	A proposed taxonomy of knowledge flows	20
3.1	Broadcasting model diffusion curve	27
3.2	Word-of-mouth model diffusion curve	28
3.3	Transition from a regular to a random network, dependent on the rewiring probability p	35
3.4	Von Neumann neighbourhood with visible range equal to 3	42
4.1	Moore neighbourhood with $v = 3$	51
4.2	Firms' Skills Universe	52
4.3	Total and joint number of innovations, by various levels of p	57
4.4	Innovation performances	58
4.5	Network statistics for good and bad run	60
4.6	Innovation performances	63
4.7	Density of the acquaintances network graph	65
4.8	Number of vertices in the largest component (acquaintances network)	66
4.9	Acquaintance and partnership network graphs: good run snapshot	68
4.10	Acquaintance network graph: good run snapshot	70
4.11	Partnership network graph: good run snapshot	72
4.12	Acquaintance and partnership network graphs: bad run snapshot	74
4.13	Acquaintance network graph: bad run snapshot	76
4.14	Partnership network graph: bad run snapshot	78
4.15	Impact of initial conditions over innovative performances	80
5.1	Knowledge diffusion patterns in the Colchagua Valley	94
5.2	Network firms–firms and network institutions–firms	98
6.1	Showing target and conceptual model	108
6.2	Validity relationships	112
6.3	Comparison for validation	114
6.4	Showing validation by comparison of models	119

7.1	Firms' geographical distribution	133
7.2	Total and joint number of innovations	136
7.3	Total number of innovations by various levels of initial knowledge and extra links	140
7.4	Single run performances around the average value	142
7.5	Innovation performances (good run vs bad run)	145
7.6	Partnership network graph, steady state configuration	148

Tables

4.1	Correlation coefficients	58
7.1	Firms' knowledge base	134
7.2	Summary statistics for the acquaintances and partnership networks	137
7.3	System performance when increasing the network density	138
7.4	System performance when increasing the network density and the knowledge base	139
7.5	Average number of innovations and standard deviation for two comparable outcomes obtained with different policy mixes	141
7.6	Total number of innovations achieved using different rules to activate new links	144
7.7	Summary statistics for the acquaintances and partnership networks (good run vs bad run)	146

PART I

Theory

1. Introduction

WHY A BOOK ON KNOWLEDGE DIFFUSION AND INNOVATION?

Economists, scientists, policy-makers and, more and more often, common people refer to modern economies as knowledge-based because of the growing relevance knowledge is acquiring in everyday life. Indeed, institutions, firms and individuals progressively rely on knowledge as a key component for individual and collective growth. This calls for a clear understanding of knowledge and its sharing patterns.

While attempting to define knowledge and investigating the complex process which determines its sharing patterns, we agree with Grant's (1996) concern that these are long-standing questions which have intrigued some of the world's greatest thinkers from Plato to Popper without the emergence of a clear consensus. Hence, in this book the focus of the investigation is restricted to the type of knowledge used by firms in the production process and, more importantly, in innovative activities.

A firm's ability to innovate depends largely upon its ability to capture and nurture human intellectual capital effectively. One important part of this process is research and development (R&D), which represents a fundamental activity for creating new knowledge for production and innovation. However, the simultaneous ongoing processes of knowledge deepening and knowledge widening – which leads to a general expansion of the range of available technologies, as well as to a growing specialization of competencies – calls for new, interactive patterns of learning.

Individual learning activities – as they are conceived in an R&D laboratory – are no longer sufficient to put together all the required knowledge it takes to be competitive. Innovative firms need specialized knowledge, as well as more types of knowledge, which increasingly lie outside the firm itself. However, because of its tacit

component,¹ knowledge, and especially new knowledge, can be difficult to acquire in the market, so firms seek some form of collaboration with other firms and/or institutions that possess the required knowledge and, on a reciprocal basis, are keen on sharing it. Hence, firms act to create links through which to access disparate and specialized resources of knowledge needed to innovate. The emerging configuration and reconfiguration of social networks of all types should then reflect the shifting demand of the knowledge economy.

This ongoing process makes it increasingly relevant to investigate the dynamics through which firms share knowledge, and calls for a thorough understanding of knowledge diffusion patterns. This entails understanding the processes through which external-to-the-firm knowledge is acquired and integrated with internal knowledge, a process which might turn out to be complex and hard to manage.

WHY THIS BOOK ON KNOWLEDGE DIFFUSION AND INNOVATION?

Now that we have explained the need for a book on knowledge diffusion and innovation, we should clarify how this book should serve the purpose of bridging the gaps in the existing understanding of knowledge diffusion and innovation. In the field of knowledge-related studies complexity arises at several levels. First, knowledge should be understood as a complex system which goes well beyond the dichotomous nature of information. Acquiring knowledge, from whatever sources, entails cognition and complex integration processes: as pointed out by Ancori et al. (2000), the economics of knowledge differs from the economics of information in the sense that knowledge is no longer assimilated to the accumulation of information in a stockpile. The distinction between these two concepts has been repeatedly ignored by a certain branch of the economic literature (economics of information), which does not consider the cognitive structure that agents use to elaborate knowledge.

Following this distinction, Ancori et al. developed a theory in which knowledge is acquired 'by a backward process through which the new knowledge is confronted and articulated with previous experience . . . The appropriation of crude knowledge – i.e. its integration in one's cognitive context – is not the result of a transmission, but rather the result of a re-engineering process' (2000, p. 267). Hence

knowledge is a complex phenomenon which requires a complex and costly cognitive process in order to be acquired. However, knowledge diffusion is not the only possible way of sharing competences. For instance firms can pool together their specialized knowledge on specific projects. Such a knowledge integration mechanism does not entail knowledge transfer.

This being said, the main elements of novelty of this study rest precisely on the complex approach undertaken to study the phenomenon under investigation. In this book we provide a definition of knowledge which is grounded in recent studies on complexity theory and, subsequently, use an agent-based social simulation methodology to address the issue of innovation – as we believe that there is great potential in addressing studies on complex social systems employing agent-based simulation models. In areas dominated by complex phenomena (such as modelling social systems) agent-based models represent, in the authors' view, a new and promising tool for scientific computational studies.

THE BOOK STRUCTURE

The book is structured in two parts. In the first part the existing literature on knowledge economics is reviewed and the issues of knowledge complexity, and knowledge and innovation, are introduced. Specifically, we first review the main literature on the knowledge-based economy, focusing on the important link between knowledge and innovation. We focus our attention on various definitions of knowledge (distinguishing between knowledge and information, as well as between tacit and codified knowledge), on the relevance of the geographical dimension for knowledge diffusion, and finally on various patterns of diffusion associated with knowledge flows (distinguishing among various forms of voluntary and involuntary-based knowledge-sharing patterns) (Chapter 2).

Subsequently, we introduce the issue of modelling knowledge and its sharing patterns. We depart from classical studies on social learning, where the patterns of information and knowledge diffusion are explored with respect to innovation adoption dynamics, and proceed to review more recent models where knowledge is considered and modelled as a complex concept (Chapter 3).

This literature review leads us to the core idea of the book, that

knowledge, and the learning processes associated with it, needs to be modelled using complex representations and appropriate tools. Critical factors in formal modelling concern the representation of knowledge (for example whether as a scalar or as a vector), the characteristics of the network structure upon which knowledge interactions (and innovation) take place, and also the temporal aspects of knowledge diffusion – simulations being sensitive to initial conditions and to the application of specific updating mechanisms.

All these factors are explored in Chapter 4 of the book, where we present an original agent-based model of knowledge diffusion, grounded in complex definitions of knowledge and network relations. In addition, the diffusion model is related to innovation processes where innovation stems from the recombination or integration of knowledge by means of a cognitive process which could be conducted either individually or collectively.

The second part of the book is dedicated to applications and empirical studies. This part opens with a chapter (Chapter 5) in which several empirical studies on the measurement of knowledge flows are reviewed. Subsequently, Chapter 6 presents a methodological investigation which first examines two alternative ways of doing research with agent-based modelling. These are theoretical and applied studies, incorporating agent-based models as a means of investigation through simulation. This is followed by a closely related discussion of validation of agent-based models. Here, validation is considered quite broadly, encompassing both inputs and outputs to the modelling as well as all stages of the model building and analysis.

In Chapter 7 an applied version of the knowledge diffusion model developed in Chapter 4 is presented. This is a calibrated model which makes use of data collected from field work conducted in Italy. The aim of this chapter is to test the validity of the model against a real-world case study, providing at the same time an exemplification of how validation of an applied model can be conducted. Chapter 8 concludes the book and presents several ideas for future research.

NOTE

1. Tacit knowledge is a type of knowledge that cannot be codified and, therefore, requires direct experience and personal interactions in order to be communicated. We will return to this concept in Chapter 2.

2. Knowledge economy: old and new issues

Due to the growing competitive pressure coming from emerging economies, modern manufacturing industries in developed countries have progressively shifted their focus from the physical processes of production to the design and marketing phases and, more relevantly, to the innovation of new products and production processes. In fact, in a globalized and competitive environment, the only viable way for firms operating in rich countries to enhance competitiveness is constantly to empower their innovative capabilities.

Innovation, defined as the process by which firms master and put into practice new product designs and manufacturing processes (Nelson and Rosenberg 1993), has to be understood as a process in which ‘new knowledge or new combinations of old knowledge are embodied in products and production processes and possibly introduced into the economy’ (Oerlemans et al. 1998, p. 4). Hence, innovation crucially involves the use of existing knowledge, as well as the ability to generate and acquire new knowledge (Howells 2002, p. 872). This view, shared by many scholars, supports the idea that firms progressively rely on knowledge as a key input of successful and long-lasting innovating activities (Pinch et al. 2003; Forsman and Solitander 2003). In other words, firms’ long-term competitiveness is highly dependent on their ability to innovate and, therefore, on their ability to enhance their knowledge base (Florida 1995; Cooke 2001; Malmberg and Maskell 2002).¹ The knowledge base of a firm could be defined as the collective character of the knowledge which depends both on individual human resources and on the mechanisms of interaction within the organization (Saviotti 1999).

The existing literature has identified at least two broad ways in which firms can enrich their knowledge base: through the use of the internal resources of the firm as well as through the use of resources located externally to the firm. On the one hand, ‘[l]earning to use

internal resources can be accomplished in several different ways, for example through R&D activities or learning by using or doing' (Oerlemans et al. 1998, pp. 3–4). On the other hand, the external mobilization of resources, which has been labelled 'learning by interacting' (Lundvall 1988, p. 362), involves the use of knowledge and experience of other economic actors (Oerlemans et al. 1998, pp. 3–4). Along these lines, David and Foray (2002) suggest that the 'gear change' in the growing speed and intensity of innovation observed over the recent decades comes about through the intensification of formal R&D activities, that is, working 'offline', isolated and 'sheltered' from the regular production of goods and services, but also, and perhaps more importantly, by learning 'online' where individuals can assess the acquired knowledge and hone their practices for future activities. Thus, understanding the sources of innovation and competitiveness in modern economies calls for a clear understanding of knowledge creation and its sharing patterns (that is, learning activities *latu sensu*).

In this chapter we will try to summarize some of the issues related to the knowledge economy. Specifically, we will focus our attention on various definitions of knowledge (distinguishing between knowledge and information, as well as between tacit and codified knowledge); on the relevance of the geographical dimension for knowledge diffusion (distinguishing between physical and relational proximity); and finally, on various patterns of diffusion associated to knowledge flows (distinguishing among various forms of voluntary- and involuntary-based knowledge-sharing patterns). Finally, we will introduce the issue of modelling knowledge and its sharing patterns. This will serve as an introduction to the analysis developed in Chapter 3.

DEFINING KNOWLEDGE

The growing knowledge flow which characterizes the so-called 'knowledge society' has made organizations increasingly concerned with the problem of selecting and organizing knowledge in a cost-efficient manner. As put by Johnson and Lundvall (1994), firms are becoming 'learning organizations' which define their internal organizational models in order to enhance their learning capabilities. The authors distinguish between different kinds of knowledge which

can be summarized as follows: (1) 'know what'; (2) 'know why'; (3) 'know how'; and (4) 'know who'.

The first type of knowledge refers to knowledge about 'facts' (for example how much does a firm invest in physical capital? how many bones compose the human skeleton? what codebook contains a specific law? and so on); it relates directly to the concept of information. As observed by Lundvall and Foray: 'there are complex areas where experts must hold a great deal of this kind of knowledge in order to fulfil their jobs – practitioners of law and medicine belong to this category' (Lundvall and Foray 1998, p. 116).

The second kind of knowledge refers to 'the scientific knowledge of principles and laws of motion in nature, in the human mind, and in society' (Lundvall and Foray 1998, p. 116). It is extremely relevant to speed up the advances in technology and to reduce frequency of errors in trial-and-error development processes. It serves as a key input for technological progress, and it is typically embedded in a skilled labour force.

The third kind of knowledge refers directly to personal skills and should be explicitly interpreted as the capability of being able to do something. Know-how has been typically associated with the kind of knowledge developed and kept within the firm. However, 'as the complexity of the knowledge base increases, a mix of division of labour and co-operation is also tending to develop in this field' (Lundvall and Foray 1998, p. 116).

The fourth kind of knowledge (know who) is similarly becoming increasingly relevant for firms nowadays. Know who refers directly to the kind of information 'about who knows what and who knows how to do what' (Lundvall and Foray 1998, p. 116). Because specialized knowledge and skills are essential elements for innovation and 'are widely dispersed due to the highly developed division of labour among organizations and experts' (Lundvall and Foray 1998, p. 116) knowledge of whom to contact is particularly valuable in this context.

Behind this classification of various kinds of knowledge there is a general distinction between knowledge and information which it is worth thinking upon. As it clearly emerges, knowledge and information are interlinked concepts; however it would be incorrect to refer to the learning activity (independently of the kind of knowledge we are referring to) simply as the accumulation of information.

Knowledge and Information

As recognized by many scholars, knowledge differs substantially from information (see, among many others, Foray 2004; Steinmueller 2002). However, Ancori et al. observed how the neo-classical approach of economics adopts a vision that ‘allows the reduction of knowledge to information, or more precisely allows knowledge to be considered a stock accumulated from interaction with an information flux’ (Ancori et al. 2000, p. 259). In a typical neoclassical fashion the universe can be described by a finite set of states to which probabilities can be assigned (Laffont 1989). In this view: ‘[k]nowledge improves when the probability of a particular state is estimated more accurately’ (Foray 2004, p. 4). Hence, knowledge is assimilated to information upon a vector of probability related to a predetermined set of states.

This view has recently come under criticism as it does not allow such important concepts as learning and cognition (Foray 2004). As maintained by several scholars, knowledge and information should be considered as two distinct concepts: the latter taking the form of structured data which can be easily transferred through physical supports, and the former involving cognition (see, for instance, Steinmueller 2002; Forero-Pineda and Salazar 2002; David and Foray 2002).

In Howells’s view:

knowledge can be defined as a dynamic framework or structure from which information can be sorted, processed and understood . . . Knowledge is therefore associated with a process that involves cognitive structures which can assimilate information and put it into a wired context, allowing actions to be undertaken from it. (Howells 2002, p. 872)

Hence, knowledge is generated through a cognitive process within which information is articulated with other information. This process, which we can label ‘learning’, allows actors to undertake actions which require the use of the acquired knowledge.

The full meaning of this distinction becomes clearer – maintains Foray – when one looks at the differences between the reproduction processes of knowledge and information: while the cost of reproducing information amounts solely to the physical cost of making a copy (for example the cost of a photocopy, the cost of duplicating

an electronic file, and so on), the cost of reproducing knowledge is much higher as it involves a cognitive process in order to disarticulate knowledge, transfer it to someone else, and rearticulate it for further use (Foray 2004). 'Knowledge reproduction has therefore long hinged on the "master-apprentice" system (where a young person's capacity is moulded by watching, listening, and imitating) or on interpersonal transactions among members of the same profession or community of practice' (Foray 2004, p. 4). Hence, reproducing knowledge involves an intellectual activity (based upon the direct interaction between the senders and the recipients), whereas reproducing information simply involves duplication.

Saviotti (1998, 1999) takes a similar stand when he says that information is of a factual nature, whereas knowledge establishes generalizations and correlations between variables. Factual information, maintains Saviotti, has nothing to do with meaning: in a firm, for instance, 'economically relevant types of information will be those required to describe the organisation itself, that is the types of human resources, capital equipment, etc., used (internal information), or those required to describe the external environment of the firm (external information)' (Saviotti, 1999, p. 126). On the contrary, in Saviotti's view, knowledge should be seen as a correlational structure: each piece of knowledge establishes correlations over some variables and over particular ranges of their values. This implies that knowledge has a local character, the degree of which can be measured by the number of variables involved and by the degree of correlation linking these variables. Given the local nature of knowledge, this implies that some information can be understood only in the context of a given type of knowledge: 'For example, the condition to be used and the sequence of operations to be followed in order to prepare a particular composite material can only be understood by someone who knows some macromolecular chemistry and physics.' (Saviotti 1999, p. 126). Hence, knowledge plays a crucial role in understanding and using information, which without sufficient prior knowledge, remains useless or inadequate for any economic application.²

Tacit and Codified Knowledge

After having assessed the existence of a clear distinction between information and knowledge, we can further elaborate on the definition of

knowledge itself. As mentioned above, knowledge has to be articulated in order to be transferred. This is because knowledge is, in its original form, completely embedded in the mind of the person who first developed it. In other words, we could say that knowledge is originally created as tacit: '[t]ypically, a piece of knowledge initially appears as purely tacit – a person has an idea' (Cowan and Foray, 1997, p. 595); subsequently this tacit knowledge can be codified by means of a cognitive process which involves its articulation.

Before reasoning on the codification process, we need to clarify better what tacit knowledge is. Rosenberg defines tacit knowledge as: 'the knowledge of techniques, methods and designs that work in certain ways and with certain consequences, even when one cannot explain exactly why' (Rosenberg 1982, p. 142). A similar definition is provided, in the same year, by Nelson and Winter in their seminal work on organizational routines and technological change. According to the authors: '[t]he knowledge that underlines skilful performance is in large measure tacit knowledge, in the sense that the performer is not fully aware of the details of the performance and finds it difficult or impossible to articulate a full account of those details' (Nelson and Winter, 1982, p. 73). This does not imply, for some authors, that tacit knowledge and skills are the same. As observed by Senker: '[s]kill implies knowing how to make something happen; it involves cognition but also other aspects such as manual dexterity or sensory ability' (Senker 1993, pp. 209–10). Whereas the acquisition of tacit knowledge is exclusively a perceptual, cognitive process.

Note that these definitions do not differ from the original idea of tacit knowledge introduced by Polanyi (1967). The tacit dimension of knowledge corresponds, in the view of the Hungarian polymath, to the form or component of human knowledge distinct from, but complementary to, the knowledge explicit in conscious cognitive processes. In Polanyi's view, we know more than we can tell, where the portion of knowledge possessed and not communicable is the essence of tacitness.³

However, in different moments in time and across different individuals, a different proportion of knowledge will be tacit and a different proportion will be codified. Hence, tacitness is a contextual rather than an absolute situation, this depending explicitly on the process of codification to which we shall now refer. Saviotti observes that: 'codification amounts to the gradual convergence of the scientific community and of other users on common concepts and

definitions (standardization) and on common contents and theories (selection)' (Saviotti 1998, p. 848). Similarly, Cowan and Foray noted how 'as the new knowledge ages, it goes through a process whereby it becomes more codified. As it is explored, used and better understood . . . more of it is transformed into some systematic form that can be communicated at low cost' (Cowan and Foray 1997, p. 595). The relevance of codification for economic purposes has been widely debated. The core argument put forward is that codified knowledge, when compared to tacit, can be transferred more easily, more quickly and at lower costs. Cowan et al. (2000) argued in favour of codification stating that uncodifiable (unarticulable) knowledge is not very interesting for social science. This stance is criticized by Johnson et al. (2002) who contest the view that codification always represents progress. According to these authors, tacit knowledge is a relevant component in human training, including the kind of training provided in institutions such as schools, universities and research institutes: '[I]f all important knowledge was in a codified form, training arguably could rely on abstract modelling, and the direct face-to-face interaction could be substituted by e-learning and electronic networks connected to external users of knowledge' (Johnson et al. 2002, p. 249).

This statement of Johnson et al. introduces a key point for us in the debate: tacit and codified knowledge flow in very different ways. Specifically, once codified, knowledge can be stored in a mechanical or technological way, as in manuals, textbooks or digital supports; it can be transferred from one person to another relatively easily, incurring the effort of getting access to the source of codified knowledge and decoding it for further use. In this respect, as observed by Steinmueller (2000), the context and intended recipient of the decoded knowledge makes a great deal of difference to the costs and feasibility of the initial codification.⁴ However, if appropriately codified (that is, codified keeping in mind the intended recipient), knowledge can be easily transferred, also taking considerable advantage of modern information and communication technologies. On the contrary, '[d]ifferent methods like apprenticeship, direct interaction, networking and action learning that include face-to-face social interaction and practical experiences are more suitable for supporting the sharing of tacit knowledge' (Haldin-Herrgard 2000).

Haldin-Herrgard identifies five main difficulties associated with tacit knowledge flows, related to perception, time, value, language

and distance. Perception refers to the characteristic of unconsciousness which entails a problem of people not being aware of the full range of their knowledge (Polanyi 1958). The time issue refers to the fact that the internalization of tacit knowledge takes a long time as it involves direct experience and reflection on these experiences. Value is contentious as many forms of tacit knowledge, such as intuition and rule of thumb, have not been considered valuable, lacking the status of 'indisputable methods'. Difficulties with language lie in the fact that tacit knowledge is held in a non-verbal form and hence involves extra effort to be shared; moreover, it has been observed that its context-specific nature makes it spatially sticky, since two parties can only exchange such knowledge effectively if they share a common social context, and thus important elements of this social context are defined locally (Gertler, 2003, p. 79). Finally, the issue of distance relates to the need for face-to-face interaction for the diffusion of tacit knowledge: since tacit knowledge is best acquired experientially it is with difficulty exchanged over long distances (Gertler 2003, p. 79). These difficulties pave the way to what could be labelled the 'geography of knowledge' debate which will be addressed in the next section.

KNOWLEDGE AND GEOGRAPHY

As already mentioned, modern economies must cope with a growing globalization process which changes the competitive environment substantially, requiring firms to redesign their competitive advantage and reposition themselves in globalized markets. In light of these changes some researchers have argued that globalization renders the significance of location for economic activity increasingly irrelevant (O'Brien 1992; Cairncross 1997). However, this opinion is not shared by many scholars who, from different points of view, argue that globalization actually increases rather than reduces the importance of location, that it endorses economic uniqueness, and that local clusters become increasingly relevant in the promotion and diffusion of innovations (Krugman 1996; Porter 1998; Fujita et al. 1999). In the words of Porter (1998, p. 90):

in a global economy – which boasts rapid transportation, high speed communications and accessible markets – one would expect location

to diminish in importance. But the opposite is true. The enduring competitive advantages in a global economy are often heavily localized, arising from concentrations of highly specialized skills and knowledge, institutions, rivalry, related businesses, and sophisticated customers.

Following this line of reasoning we shall argue, along with many scholars, that tacit knowledge is a key determinant of the geographical concentration of innovative activity, since its prominent role in the process of learning by interacting tends to reinforce interpersonal and localized relations. At the same time innovation is increasingly based on the interactions and knowledge flows between economic entities such as firms (customers, suppliers, competitors), research organizations (universities, other public and private research institutions), and public agencies (technology transfer centres, development agencies), which occur mainly locally (Gertler 2003, p. 79). For a growing number of scholars, this explains both the perpetuation and deepening of geographical concentration in a world of expanding markets (Gertler 2003, p. 76), as well as the spatial concentration of research and development activities in the home base of the innovating firms – defined as an important case of non-globalization (Pavitt and Patel 1991).

The question at stake is why tacit knowledge remains spatially sticky in a globalized world characterized by increasingly cheaper and pervasive diffusion of communication technologies. A first answer relates to the very nature of tacit knowledge. As discussed above, several authors maintain that knowledge in its tacit form can be effectively exchanged only by means of direct face-to-face encounters. In a nutshell, tacit knowledge could be defined as ‘person-embodied, context-dependent, spatially sticky and socially accessible only through direct physical interactions’ (Morgan 2004, p. 12). Some scholars maintain that these features of tacit knowledge help to explain the apparent paradox of phenomena such as the ‘economically successful industrial clusters in an age in which new telecommunication systems facilitate the transfer of ever more complex sets of knowledge at an ever-increasing rate’ (Pinch et al. 2003, p. 375).

Nonetheless, the debate over present and future trends is still open. As observed by Haldin-Herrgard (2000), high technology facilitates tacit knowledge diffusion in artificial face-to-face interaction, through different forms of meetings in real time, using, for

instance, audio and video conferences. This perspective is shared by other scholars; in a recent paper Brökel and Binder stated, for instance, that: '[n]ew information technologies, for example, video conferences, cast doubt on the advantages of face-to-face contacts' (Brökel and Binder 2007, p. 154).

A second, and perhaps stronger, answer relates to the importance of trust among subjects exchanging knowledge if genuine learning is to occur (Collins 2001; Nonaka and Takeuchi 1995). Scholars emphasize that in interactions amongst firms, trust has a two-way nature that can be considered a relational asset. It is more likely to develop where the participants are engaged in several encounters, meaning that 'the shadow of the future' looms larger over the present (Axelrod 1984; Morgan 2004). 'This provides a context for reciprocity: a good deal for informal know-how trading takes place, even among rival firms, precisely because of the expectation of the information which A provides B today will be reciprocated in kind tomorrow' (Morgan 2004, p. 8). This idea, first introduced by von Hippel (1987), leads to the conclusion that mutual trust and reciprocity are easier to sustain in the context of geographical proximity (Malmberg 1997).

From this perspective, tacit knowledge is considered to be context-dependent: 'being facilitated by a common language, cultural and value system. Codifiable knowledge, by contrast, can be expressed in various forms, and rapidly disseminated through various geographically dispersed user communities' (Pinch et al. 2003, p. 375). Along this line of reasoning, Maskell and Malmberg (1999) coin the concept of 'ubiquitification' of knowledge. By means of codification, 'many previously localized capabilities and production factors become ubiquities. What is not ubiquified, however, is the non-tradable/non-codified result of knowledge creation – the embedded tacit knowledge that at a given time can only be produced in practice' (Maskell and Malmberg 1999, p. 172).

This recontextualization of tacit knowledge opens up to a new specification of proximity. If the diffusion of tacit knowledge is facilitated by a 'shared language' (Burns and Stalker 1961) or, as put by Dosi and Marengo (1994), by a 'shared cognitive framework', then there is room to believe that tacit knowledge dissemination is subject to organizational or relational proximity more than to physical or geographical proximity. In a similar fashion Breschi and Lissoni (2003) refer to this kind of proximity as 'social distance': a tight

social connection facilitates learning among subjects. Hence, knowledge diffusion is not only a geographically spatial phenomenon, it is also a socially spatial phenomenon. So space matters for knowledge diffusion, and social space may matter as much as or more than geographic space (Cowan 2004, p. 3).

However, the view which juxtaposes relational proximity on the one hand with geographical proximity on the other has come under criticism of being:

a form of *spatial fetishism* which . . . lies in the assumption that there is something called *geographical proximity* which does not involve *relational proximity*, implying that the social interactions which constitute *local* actions are somehow natural, primordial or automatic, when in fact they have to be actively constructed like any other relational asset, whatever the spatial scale. (Morgan 2004, p. 11)

In other words, relational proximity, far from being a substitute for geographical proximity, should be considered as a complementary and reinforcing element of it.

This idea lies at the heart of the learning region approach. According to the regional economic geography, planning and innovation systems literature (see Florida 1995; Asheim 1996; Morgan 1997; Cooke and Morgan 2000; Maskell and Malmberg 1999; Lundvall and Maskell 2000), tacit knowledge is not well suited to flow between distant spatial locations. It is argued that face-to-face, local interaction remains important for establishing trust between partners, enabling mutual understanding, and more effectively facilitates flows of tacit knowledge.

KNOWLEDGE FLOW PATTERNS

As it emerges from the discussion conducted so far, the distinction between tacit and codified knowledge lies at the very heart of the problem of understanding knowledge flow and innovation patterns. However, in our view, the existing literature has neglected to classify the different ways in which knowledge can flow between agents. This has created some confusion and has generated a misuse of specific concepts. To help clarify the material reviewed in this section, following Morone and Taylor (2008), we present a taxonomy of knowledge flows classifying the different forms of sharing patterns.

Knowledge Gain vs Knowledge Diffusion

We start our analysis by distinguishing between the two broad concepts of knowledge gain and knowledge diffusion. The first relates, in our view, solely to those processes of knowledge flows which deliberately involve a barter among subjects: a portion of subject A's knowledge flows to subject B, who pays her back either with a portion of her knowledge or in a different way.

We shall refer to the first of these two options (that is, knowledge is paid back with other knowledge) as knowledge exchange, and to the second option (that is, knowledge is paid back with a different coin) as knowledge trade. An example of knowledge exchange has been used by Cowan and Jonard who define a model in which knowledge flows 'through barter exchange among pairs of agents' (Cowan and Jonard 2004, p. 1558).⁵ Patterns of knowledge trade, on the other hand, relate for instance to those cases where disembodied knowledge flows through technology and patent trades (Arora et al. 2002).

Note that knowledge gain relates to both tacit and codified knowledge. Codified knowledge can flow among distant agents, whereas tacit knowledge gains require always a direct interaction (that is, face-to-face) between agents.

Substantially different is the concept of knowledge diffusion. Here knowledge is no longer traded on a voluntary basis (*quid pro quo*), but freely flows while agents interact. Several scholars have referred to this process as knowledge spillover (Jaffe 1986), or knowledge percolation (Antonelli 1996). The common idea behind these definitions is that knowledge flows freely, within a specific space, and can be economically exploited by the recipient agent. The kind of knowledge being spilled over is tacit in nature, and requires some 'absorptive capacity' to be effectively recombined in the cognitive framework of the recipient agent.⁶

This definition, however, is still vague as it does not really describe how knowledge actually flows across agents, nor how it is recombined with the existing knowledge of the recipient. We shall suggest a decomposition of the knowledge diffusion concept into three sub-categories: knowledge spillover, knowledge transfer and knowledge integration.

Decomposing Knowledge Diffusion

Knowledge integration refers to a process which combines dispersed bits of knowledge held by individuals to be applied in a coordinated way, and only on a temporary basis. On the contrary, knowledge spillover and knowledge transfer denote two similar processes in which bits of knowledge conveyed from one agent to another are such that the recipient can absorb them into her/his already existing personal knowledge (that is, there is learning, hence some previously acquired related knowledge is required); the only difference between these two processes being that spillovers are unintended processes of knowledge diffusion (for example while chatting with colleagues) whereas knowledge transfer requires a defined will (for example while jointly working on a project).⁷

Knowledge transfers and knowledge spillovers are the most cited typologies of knowledge diffusion patterns (see, for instance, Morone and Taylor 2004a; van der Bij et al. 2003; Cabrera and Cabrera 2002; Hansen 1999; Szulanski 1996). However, these mechanisms present some disadvantages: they are expensive and often time-consuming and they offset the specialization of employees needed for innovation as they both assume that individuals absorb diverse specialized knowledge by means of face-to-face encounters (Demsetz 1991). In fact, here we are posing a question of depth of knowledge versus breadth of knowledge; as suggested by Grant: '[d]ue to cognitive limits of the human brain, knowledge is acquired in a highly specialized form . . . However, production – the creation of value through transforming input into output – requires a wide array of knowledge, usually through combining the specialized knowledge of a number of individuals' (Grant 1996, p. 377). The possibility to integrate knowledge without having to acquire it might provide a solution to these drawbacks. In light of these arguments, Grant asserts that integration of specialized knowledge is at the heart of production in a knowledge-based society.

But how does integration occur? In a recent paper Berends et al. (2006) examined knowledge integration in industrial context and elaborated the concept of 'thinking along', that is, a mechanism that allows for knowledge integration without the need for transfer. More precisely, thinking along is the process whereby an agent applies knowledge temporarily to a problem of somebody else and communicates the generated ideas to that other person. Hence, it




Figure 2.1 A proposed taxonomy of knowledge flows

involves a temporary cognitive work with regard to a problem of someone else.

Interestingly, the concept of knowledge integration does not involve any permanent flow of knowledge from subject A to subject B in the conventional sense. We consider it an ‘atypical’ form of knowledge diffusion. Knowledge integration is an important catalyst for knowledge flow – though unpredictable – generating insights and unsticking the workflow. Moreover it is itself a form of meta-knowledge (that is, knowledge about knowledge) flow.

Now, recombining the analysis developed in this section, we shall propose a taxonomy of knowledge flow processes. Figure 2.1 shows a taxonomy of concepts emerging from an analysis of the knowledge flows literature. At the top level in the hierarchy are knowledge gain and knowledge diffusion which we classify as distinct phenomena of flows. Knowledge exchange and trade are subclasses of knowledge gain, whereas knowledge spillover, transfer and integration are derived from a decomposition of knowledge diffusion.

Not only does this analysis provide a background for comparing studies of knowledge flows, but in doing so it highlights different assumptions about governance and control of knowledge. In the case of knowledge gain one assumes the functionality, as well as the ability, of locking flows in a rigidly controlled domain of knowledge. The strategy is to maximize the pay-off of current knowledge assets

and obtain a fair value in exchange. The drawback of this approach is that over the long term it tends to stifle creativity and diminish diversity in production of new knowledge and recombination of existing knowledge. The opposite strategy is the promotion of largely uncontrolled diffusion, where value is often derived from the outcomes on a larger scale: the generation and exploitation of whole new economic areas, and the impact this has on the opportunities and constraints for the organization.

Indeed, in discussing the flows between individual units of the knowledge economy it is important not to forget macroeconomic features arising from the specific structures. Knowledge dynamics can be conceived of as systemic outcomes and patterns based on the interplay and aggregation of flows, and therefore involving scaling up towards large systems of interacting agents and across longer temporal frames. Specific flows can be used in explaining the performance of the knowledge economy, its distributional outcomes (for example equality), innovation and adoption patterns (of codification supports for example).

CONCLUSIONS AND NEW CHALLENGES FOR RESEARCHERS

The main objective of this chapter has been to introduce the reader to several issues related to knowledge economy. We have done so by presenting a review of recent studies on knowledge, arriving at some definitions of terms in the knowledge economy field. As it has emerged, knowledge is indeed a key input for innovation; however, from a theoretical point of view, much effort is needed to capture its role in innovating activities due to the complex nature of learning processes.

In the area of knowledge flows several different types of modelling have been used. Conceptual modelling ranging from organizational models (often used to focus on company activities) to taxonomic models (such as the one presented above) are found. Mathematical modelling can be used to determine solution states and optimization behaviours. On the other hand, simulations are promising tools with which to investigate knowledge flows because they can express the dynamics in a model. There is no standardized way of modelling, despite the nowadays widespread use of computers and simulation.

Computational approaches can be quite complicated, involving many model components and parameters that need to be calibrated. Various aids – mathematical notation, set theory, logical forms, flow diagrams, decision trees – are often used to stage between a verbal description of modelling concepts and the computer code of the implementation.

In the next chapter we will present a survey on models of knowledge flows, departing from earlier works on new technology adoption (the so-called epidemic models), reaching later studies which make use of game-theoretical tools, and concluding by reviewing increasingly complex simulation models.

The role of formal modelling and simulation is to allow exploration of the hypotheses embodied in the program over a range of different conditions. As we will discuss in the second part of this book, models can be, to a greater or lesser degree, based on empirical data on knowledge flows. Although measurement is often problematic (we will discuss the measurement issue in Chapter 5), efforts to improve the empirical basis of modelling are key to the increasing sophistication of recent knowledge diffusion models. When this improvement or ‘model validation’ step is done as part of iteration with model design and simulation experiment, it can open the opportunity for dialogue between modeller and empirical analyst or expert. This leads to describe a second – and equally important – role of modelling: to improve the clarity of conceptual models and help the modeller to arrive at a more rigorous conceptualization. We shall address the issue of model validation in Chapter 6, where a methodological discussion on applied and theoretical modelling is presented. Finally, in Chapter 7 we will present a model validation exercise.

NOTES

1. We will get back to the link between knowledge and innovation in Chapter 3 and Chapter 4 where we will first present a survey on the existing literature, and subsequently an agent-based simulation model.
2. We revisit the distinction between knowledge and information in the next chapter, where we will address the issue of knowledge as a complex phenomenon.
3. It should be mentioned that knowledge is rarely completely tacit or completely codified. In most cases, a piece of knowledge can be placed somewhere in a range that stretches from the completely tacit to the completely codified.
4. To grasp this concept, one may think of the difference in codifying economic

notions for undergraduate and postgraduate students. Indeed, given the different intended recipients, the same notions will be codified in rather different ways, entailing different kinds of effort.

5. We will come back to Cowan and Jonard's (2004) model of knowledge exchange in Chapter 3.
6. Note that the diffusion process here defined means that knowledge is a non-rival and relational asset (that is, an asset the use of which by one agent does not prevent a simultaneous use by other agents and the diffusion of which is based on reiterated personal interactions). These characteristics have been widely discussed in the literature (see, among many others, Nelson 1959; Arrow 1962; Axelrod 1984; von Hippel 1987). Some authors have referred to this as knowledge being a club good (see, for instance, Breschi and Lissoni 2001).
7. On this issue see Brökel and Binder (2007). According to the authors: 'Intended knowledge transfers are when actors actively seek knowledge ('search'). Unintended knowledge transfers might be considered as an individual "stumbling upon" knowledge, for example, while visiting a trade fair or listening to a presentation' (2007, p. 155).

3. Modelling knowledge and its diffusion patterns: a pathway towards complexity

The basic assumptions of most neoclassical economic models are perfect information and homogeneous technology. These imply that all economic agents (for example producers) possess at any moment in time the same identical information upon available technologies and produce with the same production function.¹ From these hypotheses it follows that information, which in this context encompasses knowledge, is a public good (that is, non-rival and non-excludable) freely available to all in the economy. However, these assumptions have failed empirical tests. Jaffe et al. (1993) showed quite clearly that knowledge spillovers (measured by patents citations) are geographically localized and that geographic localization fades over time very slowly. This implies that the diffusion process of knowledge follows specific patterns and is not at all an automatic and instantaneous process.²

In this chapter several knowledge diffusion models are presented, commencing with models in which the diffused knowledge is equated with information about new technology (that is, innovation), and proceeding to review more recent models where knowledge is considered and modelled as a complex concept.

REVIEWING INNOVATION DIFFUSION MODELS

As discussed in Chapter 2, most of the learning processes occur among neighbouring agents, given that knowledge is, at least to some extent, persistently and inherently local in nature. Pioneering studies on diffusion investigated the patterns through which new technologies are adopted and spread in social systems (see, for instance, Rogers

2003 [1962]; Hägerstrand 1967; Casetti and Semple 1969; Bernhardt and MacKenzie 1972; Sahal 1981). Such investigations could be considered the first studies on knowledge diffusion, considering the underlying hypothesis that a new technology diffuses when sufficient information on the characteristics of the new technology spreads from earlier adopters to later adopters. Hence, the technology diffusion process resembles the underlying knowledge/information diffusion dynamic. Reversing the argument, since new technologies can be adopted when sufficient information is available: 'one is likely to learn a lot about the time path of technology diffusion by studying the spread of information about it' (Geroski 2000, p. 604).

However, some empirical studies found a lag of several years between the date when farmers learned of the existence of new technologies and the date when they adopted it. Slicher van Bath (1963, p. 243) observed that: 'land tilled in very ancient ways lay next to fields in which crop rotations were followed'. In other words, as put by Geroski: '[s]ometimes it seems to take an amazingly long period of time for new technologies to be adopted by those who seem most likely to benefit from their use' (Geroski 2000, p. 604). These findings suggest that technology diffusion is not just a process of spreading news, but rather it involves persuasion.³

Rogers, in his seminal study on diffusion of innovations, proposed an interpretative model of innovation diffusion articulated into five stages: knowledge, persuasion, decision, implementation and confirmation. Consequently, innovators follow a specific pathway towards adoption: (1) a preliminary need for acquisition of adequate knowledge of an innovation; (2) which will help them in forming an attitude toward the innovation; and (3) decide whether to adopt or reject it. Hence, adopters will (4) put the new idea into use; and (5) seek reinforcement of the innovation decision already made; however, in this last phase adopters may reverse this previous decision if exposed to conflicting messages about the innovation (Rogers 2003 [1962], p. 169). Therefore, knowledge spread is a necessary condition in Rogers's view; yet it is not sufficient to determine the diffusion of an innovation.

Epidemic Models

A class of technology diffusion models are the 'epidemic models'.⁴ We shall distinguish between broadcasting and word-of-mouth

epidemic models. The first subclass refers to a model where the source of knowledge upon the existence and/or characteristics of the new technology is external and reaches all possible adopters the same way (broadcasting or external-influence diffusion model); whereas the second subclass refers to a model in which knowledge is diffused by means of personal interactions (word-of-mouth or internal-influence diffusion model).⁵

Let us define a population of \bar{N} potential adopters of the new technology. Each potential adopter will switch to the new technology as soon as he/she hears about its existence (hence, technology diffusion overlaps completely with information diffusion). We shall define $g(t)$ as the coefficient of diffusion (that is, the probability that at time t a new subject will adopt the new technology), which might or might not be a function of the number of previous adopters.

In the broadcasting model $g(t)$ is set equal to a constant parameter: $g(t) = a$ with $(0 < a < 1)$, which can be considered the strength of the broadcasting signal or, alternatively, the infection probability. Let us define $N(t)$ as the cumulative number of adopters at time t . The increase in adopters for each period will be equal to the probability of being infected multiplied by the current population of non-users (that is, the number of agents who are still not using the new technology). Formally we can write the rate of diffusion at time t as:

$$\frac{dN(t)}{dt} = a[\bar{N} - N(t)] \quad (3.1)$$

which, integrated, gives the cumulative adopter distribution:

$$N(t) = \bar{N}[1 - e^{-at}] \quad (3.2)$$

Equation (3.2) is a modified exponential function (with negative exponent) of the type described in Figure 3.1;⁶ the smaller the value of a the slower the diffusion process will be and the smaller will be the number of users at any time. Note that this particular information diffusion process is based on the assumption that: ‘the rate of diffusion at time t is dependent only on the potential number of adopters present in the social system at time t . In other words, the model does not attribute any diffusion to interaction between prior adopters and potential adopters’ (Mahajan and Peterson 1985, pp. 16–17). As observed by Geroski: ‘[t]his kind of model of information




Figure 3.1 *Broadcasting model diffusion curve*

diffusion is not an implausible story of how people become aware of a new yoghurt product or news about the fall of the Berlin Wall. However, technology adoption often takes an order of magnitude longer than it takes for information to spread' (Geroski 2000, p. 605). At the heart of the problem is the distinction discussed in Chapter 2 between knowledge and information. Although information can be transmitted impersonally through a users' manual, much of the knowledge required to use any specific technology proficiently is built up from the experience of using it, and at least some of that valuable knowledge will be tacit. As a consequence, it must be transmitted from person to person, and cannot be effectively broadcast from a common source (Geroski 2000, p. 605). Figure 3.1 illustrates a broadcasting model diffusion curve.

This distinction leads to the development of a new model which falls into the second subclass of epidemic models – that is, those based on word-of-mouth diffusion – and which represents a first step towards a more accurate definition of knowledge.⁷ As observed above, the basic assumption of this model is that knowledge diffuses among agents by means of face-to-face interactions. It follows that the probability of

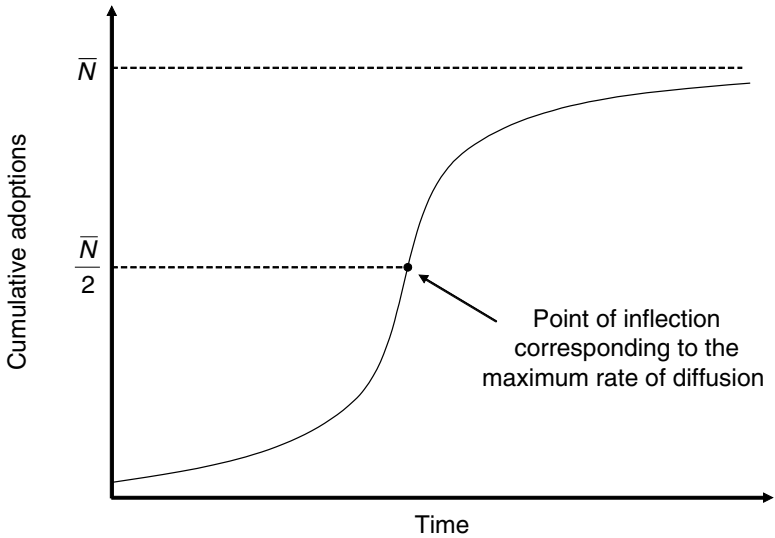


Figure 3.2 Word-of-mouth model diffusion curve

receiving the relevant knowledge needed to adopt the new technology is a positive function of current users $N(t)$. Let us define the coefficient of diffusion $g(t)$ as $bN(t)$ with $b > 0$. Hence, the fundamental equation of the rate of diffusion at time t can now be defined as:

$$\frac{dN(t)}{dt} = bN(t)[\bar{N} - N(t)] \quad (3.3)$$

This function can be integrated in order to obtain the cumulative adopter distribution function:

$$N(t) = \frac{\bar{N}}{1 + \left(\frac{\bar{N} - N_0}{N_0}\right)^{-b\bar{N}(t-t_0)}} \quad (3.4)$$

where $N(t = t_0) = N_0$.

The logistic function reported in equation (3.4) traces an S-shaped curve (see Figure 3.2) which is consistent with the dominant stylized facts. It captures the different speeds of the diffusion process: 'diffusion rates first rise and then fall over time, leading to a period

of relatively rapid adoption sandwiched between an early period of slow take up and a late period of slow approach to satiation' (Geroski 2000, p. 604). This means that while usage increases year by year over time, it does so more rapidly in the early years after the introduction of a new technology than it does after the technology has become fairly well established (Geroski 2000, p. 606).

A limitation of this second subclass of models is that it cannot explain the diffusion of an innovation from the date it is invented, but only from the date when some number, $N(t) > 0$, of early users have begun using it. In fact, as observed by Geroski: 'word of mouth diffusion processes can only begin to happen after an initial base of users has been built up, and, needless to say, the larger is this initial base of users, the faster is diffusion' (Geroski 2000, p. 606). A possible way to overcome this problem would be to integrate equations (3.1) and (3.3), obtaining a mixed information source model. In such model existing non users are subject to two sources of information, and the probability of being informed (that is, the coefficient of diffusion) is $[a + bN(t)]$; the resulting equation of the rate of diffusion at time t will be:

$$\frac{dN(t)}{dt} = (a + bN(t))[\bar{N} - N(t)] \quad (3.5)$$

The cumulative adoption distribution results in a generalized logistic curve whose shape is determined by both a and b . This version of the epidemic model is indeed the most general and the most widely used as it can accommodate the assumption underlying the two models just discussed above (Mahajan and Peterson 1985).

Further complications of this model include the definition of two populations (hence, introducing a degree of heterogeneity among agents) which might differ for their ability to understand the new technology. Such complication can be used to mimic a process in which b declines over time or a situation in which the total pool of potential users, \bar{N} , is not fixed but increases over time (Geroski 2000, pp. 608–9). This helps to overcome a limitation of most epidemic models: namely that the infection probability depends upon exogenously determined parameters.

Game-Theoretical Models

Recently, several scholars have developed new game-theoretical models which allow for more complicated logic than that behind the

basic epidemic model, introducing several sources of heterogeneity. Ellison and Fudenberg (1993) developed a model in which agents consider the experience of neighbours in deciding which of two technologies to use. Their work is structured around two simple models of the learning environment. First, they consider a homogeneous population which reacts identically to the two technologies, with one technology being superior to the other; subsequently they introduce heterogeneous agents which consider the two technologies in different ways. In the first case the issue is whether the better technology will be adopted, while in the second case the question is whether the two technologies will be adopted by the appropriate players. In both environments agents use exogenously specified 'rules of thumb' that ignore historical data but might incorporate a tendency to use the more popular technology. Under this condition, the outcome of their work suggests that 'even very naïve learning rules can lead to quite efficient long-run social states', but adjustment can be slow when a superior technology is first introduced (Ellison and Fudenberg 1993, p. 637).

In a subsequent paper (Ellison and Fudenberg 1995), the authors focused their attention on the patterns of information exchange, studying the way in which word-of-mouth communication contributes to the aggregation of individual agents' information. They defined a non-strategic environment composed of homogeneous agents which face the decision of choosing one of two products. Their findings show how: 'despite the naïve play of individuals, this type of information flow may lead to efficient learning on the social level, and that social learning is often most efficient when communication between agents is fairly limited' (Ellison and Fudenberg 1995, p. 120).

Bala and Goyal (1995) analysed social learning in an organizational environment in which agents have a time limited life experience and heterogeneous beliefs. Departing from the original work of Kiefer (1989), they consider the case of a monopolistic firm attempting to calculate its true demand curve. 'The learning is then one in which the sequence of managers learn about the demand curve through their own actions as well as the experience of earlier managers' (Bala and Goyal 1995, p. 303). In this case again, learning from others augments the probability of converging to the set of ex-post optimal actions.

Subsequently, Bala and Goyal (1998) investigated the relation

between the structure of social networks and learning processes in a world where agents learn from currently available social information, as well as from past experiences (as opposed to the previous works of Ellison and Fudenberg). Their findings show that the structure of the neighbourhood has important implications for the likelihood of adopting new technologies, for the coexistence of different practices, and for the temporal and spatial patterns of diffusion in a society. More precisely, they show how neighbourhood structures characterized by the presence of locally independent agents (that is, agents with non-overlapping neighbourhoods) generally facilitate social learning.

A common way of modelling the mechanisms of social learning and technology diffusion makes use of evolutionary game theory. Several authors examined local interaction games in which each person's pay-off depends on the actions of his/her neighbours. Most of these studies pointed out that local interaction might result in the diffusion of personal behaviours in certain dynamic systems. In a recent work, Morris (2000) extended this finding, linking the occurrence of social learning (which he calls contagion) to some qualitative properties of the interaction system such as cohesion, neighbour growth, and uniformity. Chwe (2000) modelled social learning as dependent on both social structure and individual incentives. In this way he obtained a model that he called a 'local information game' as, he argued: 'locality is represented by information and not necessarily by payoffs . . . Local interaction games make sense for local coordination, such as keeping our street clean; for "big" coordinations such as political changes, informational locality is more appropriate' (Chwe 2000, p. 11).

Although game theory allows knowledge diffusion models to overcome most of the limitations of basic epidemic models (introducing, for instance, heterogeneous agents with heterogeneous beliefs; building social networks in order to avoid the simplifying hypothesis of complete mixing of social system members; allowing for different innovations occurring at the same time; and so on),⁸ yet, none of these studies goes beyond a binary definition of learning. Given a new technology, agents can either learn of its existence, and have the possibility of adopting it, or stay in their initial state of ignorance and not adopt it. However, several classic studies (Arrow 1962; Nelson and Phelps 1966; Rosenberg 1982) showed how learning from peers can be a much more complicated process which evolves

by incremental improvements. In order to take onboard these considerations, new methodologies need to be developed. Specifically, a higher degree of complexity needs to be incorporated into knowledge diffusion models if one wants to study the actual dynamics governing learning activities.

KNOWLEDGE FLOW MODELS

In-depth empirical studies have addressed the knowledge diffusion process revealing the whole complexity of knowledge-sharing patterns associated with innovating activities. Allen (1983) observed that under the conditions prevailing in the nineteenth century technical knowledge was released relatively freely to actual and potential competitors, allowing for cumulative advance. This process, in turn, would lead to what he called 'collective invention', probably the most effective source of innovation at the time. Hence, firms share knowledge to innovate and not, as assumed in the above-reviewed literature, knowledge on innovation. This is a crucial distinction as it suggests that economists should look at knowledge as an input of innovation which can be acquired by direct interaction among agents operating in social systems (even in competitive environments). Consequently, the focus of the analysis should shift from studying information diffusion patterns to the mechanisms governing knowledge sharing and its use for innovation purposes.

An important contribution in this direction is provided by the empirical investigation of von Hippel (1988). The author observed know-how trading in a variety of industries; such trading activity between rivals is defined as a general and significant mechanism that innovators can use to share (or avoid sharing) innovation-related costs and profits with rivals. It is essentially a pattern of informal cooperative research and development. The know-how notion used by von Hippel is rather similar to the notion of tacit knowledge discussed in Chapter 2. In the author's words, know-how is: 'held in the minds of a firm's engineers who develop its products and develop and operate its processes' (von Hippel 1988, p. 76). The way in which such informal trading occurred is well described by the author (von Hippel 1988, p. 77):⁹

[K]now-how trading behavior . . . can be characterized as an informal trading network that develops between engineers having common

professional interests. In general such trading networks appear to be formed and refined as engineers get to know each other at professional conferences and elsewhere. In the course of such contacts, an engineer builds his personal informal list of possibly useful expert contacts by making private judgments as to the areas of expertise and abilities of those he meets. Later, when Engineer A encounters a difficult product or process development problem, A activates his network by calling Engineer B – an appropriately knowledgeable contact who works at a competing (or non-competing) firm – for advice.

Engineer B makes a judgment as to the competitive value of the information A is requesting. If the information seems to him vital to his own firm's competitive position, B will not provide it. However, if it seems useful but not crucial – and if A seems to be a potentially useful and appropriately knowledgeable expert who may be of future value to B – then B will answer the request as well as he can and/or refer A to other experts. B may go to considerable lengths to help A: for example, B may run a special simulation on his firm's computer system for A. At the same time, A realizes that in asking for, and accepting, B's help, he is incurring an obligation to provide similar help to B – or to another referred by B – at some future time. No explicit accounting of favors given and received is kept, I find, but the obligation to return a favor seems strongly felt by recipients – '... a gift always looks for recompense'.

Based on this empirical evidence, scholars have recently attempted to develop a new class of theoretical models able to capture the underlying complexity of these knowledge transfer mechanisms using social networks tools to govern who interacts with whom. Such models fall in the class of agent-based simulation models and form a whole new way of looking at economic problems. As suggested by Fagiolo et al., agent-based modellers: 'reject the aprioristic commitment of new classical models to individual hyper-rationality, continuous equilibrium, and representative agents. Everything in the neoclassical world can, in principle, be known and understood' (Fagiolo et al. 2007, p. 255). On the contrary, in agent-based models the set of objects in the world (for example techniques of production, or products) is unknown, and agents must engage in an open-ended search for new objects. This distinction leads to several other differences with regards to the types of innovative learning and adaptation that are considered, 'definitions of bounded rationality, the treatment of heterogeneity amongst individual agents and the interaction between these individuals, and whether the economic system is characterised as being in equilibrium or far-from-equilibrium' (Fagiolo et al. 2007, pp. 255–6). In what follows we will review some

recent agent-based models which deal with the problem of knowledge flows. Such models will be then classified on the basis of the taxonomy developed in Chapter 2.

Modelling Knowledge Exchange through Bilateral Bartering

Cowan and Jonard (2004) develop a model in the framework of graph theory. Agents are arranged in one dimensional space, where each agent occupies one vertex and may interact with her/his k nearest neighbours on either side. Knowledge is of several types, so agents' endowments are represented by a vector. A small number of agents are 'expert' and are endowed with a high level of knowledge in at least one value of the vector. When agents meet, they exchange knowledge in a barter arrangement, swapping knowledge of different types. Bartering can only take place if one individual has superior knowledge of one type and the other individual has superior knowledge of another type.¹⁰ Knowledge is a non-rival good and can be exchanged without decreasing the level of knowledge possessed by each trader. Moreover, knowledge transfer is not complete as agents possess limited absorptive capacity.¹¹

The dynamics of the model are very simple: each period one agent is chosen randomly. At random he/she selects one of his/her direct connections, and with that agent makes all possible exchanges. Over time, agents' knowledge endowments change as they interact. The main target of this model is to investigate how aggregate knowledge levels (measured as the mean knowledge level over all agents) grow in such a social system, and how aggregate growth is affected by network architecture.

To address these questions, the model is simulated using the Watts–Strogatz (1998) algorithm to create the networks over which knowledge exchange takes place. Initially agents are allocated over a regular graph. Subsequently the graph is exogenously modified. With a certain (rewiring) probability p , any individual may break off the connection with the neighbour and reconnect to a vertex chosen uniformly at random over the entire lattice. With probability $1-p$, this leaves the edge unchanged. Note that when the rewiring probability is set equal to one the resulting graph is a network with each agent being connected to, on average, n randomly chosen agents (that is, random network). As the value of p changes, the structure of the network changes. By choosing a fairly small value for p , the




Figure 3.3 Transition from a regular to a random network, dependent on the rewiring probability p

lattice remains almost regular and highly clustered (that is, with a high degree of cliquishness). However, each long-range edge provides a short cut not only between the two vertices that this link connects, but also for their immediate neighbours, the neighbours of those neighbours, and so on. Hence, the graph has the low average path length of an almost-random graph (the concept of average path length introduced here is a measure of the efficiency of the model, giving the average number of steps required to connect each pair of vertices in the lattice). In Figure 3.3 the transition from a regular network to a random network is shown.

The simulation exercise developed by Cowan and Jonard on the basis of this model leads to some interesting results. Depending on the network structure, the model produces a spectrum of states of the world varying from a situation of high knowledge inequality and fast knowledge diffusion, to the opposed situation, more equal in terms of knowledge variance but less efficient in terms of knowledge diffusion. The small world region is one of high mean and variance. As suggested by the authors, this finding would suggest some policy tension if one considers the link between knowledge distribution and wealth distribution. However, the growing variance could be due to some scaling effect due to the overall increase in the mean knowledge level of the system. Cowan and Jonard controlled for this effect by calculating the relationship between network architecture and the equilibrium coefficient of variation of knowledge levels:

This measure tells quite a different story. Normalized dispersion of knowledge levels falls in the small world region, suggesting that what is lost in terms of increased heterogeneity as measured by the variance is

more than made-up [*sic*] by an increase in overall knowledge levels. If this is an appropriate measure of dispersion, the policy tension between efficiency and equity dissolves – both are optimized with a small world network. (Cowan and Jonard 2004, p. 1565)

Investigating the transition towards the long-run equilibrium yields some interesting results. Random networks are characterized by small local structures, on the one hand, and short paths between agents, on the other hand. However, while diffusion is fast during early stages of the history, the process is exhausted at relatively low levels of aggregate knowledge. Regular networks are also locally structured although path lengths between agents are long. Moreover, while diffusion during early stages is slower, the process itself is longer, and reaches a higher aggregate level of knowledge. Finally, small world networks enjoy the advantages of both: as they have relatively short path lengths, diffusion in early periods is relatively fast; moreover, since they are locally structured, exchange continues longer than in random worlds (Cowan 2004, p. 13).

A clear implication of the dominance of small world networks over random networks is that the average path length ‘is actually not an unequivocal measure of the performance of this structure. Diminishing the distance between members of an organization or economic system by reallocating links does not always improve performance: the architecture of links matters or, put another way, there is value to cliquishness’ (Cowan and Jonard 2004, p. 1564).

Some explanations of the dominance of small world networks may be found in the specific trading mechanism employed to exchange knowledge. Specifically, barter exchange requires always that potential traders satisfy a double coincidence of wants: ‘When this fails, potentially improving transfers do not take place. Cliquish network structures mitigate this problem, because within a clique there are typically several indirect but still short paths between any two agents’ (Cowan and Jonard 2004, pp. 12–13).

Moreover, when conducting sensitivity analysis of results to parameter specification it emerges that the small world results are affected by absorptive capacity. Redoing the experiment with different values for absorptive capacity shows that the optimal network structure becomes more and more random, that is, the optimal value of p increases with absorptive capacity. This occurs because as absorptive capacity increases, the value of cliquishness falls, and thus the relative value of short paths increases.

Although this study sets a new way of modelling knowledge flows, the following limitations can be highlighted. First, the knowledge exchange mechanism developed by Cowan and Jonard (2004) ignores costs and benefits associated to bilateral bartering. Second, in the Cowan and Jonard (2004) model there is no attempt to consider situations where the social network structure may change through agent behaviour and interaction – that is, a static network structure that governs social interactions is overimposed upon the system and is exogenously modified. A third limitation is the assumption that agents interact on the basis of complete information on the level of knowledge of their acquaintances. As a matter of fact, it is not in the spirit of the study of knowledge generation and diffusion to assume a priori what agents need to learn. In fact, as discussed above, knowledge generation and diffusion models are rooted in the Nelson and Winter evolutionary approach, which rejects the hyper-rationality approach where everything is known and clearly understood, building on the concept of bounded rationality according to which agents must engage in open-ended learning processes in order to acquire new knowledge and information. A further limitation is the adoption of an oversimplifying notion of knowledge. Considering knowledge as a number (or as a vector of numbers) is indeed a convenient simplification, but restricts our understanding of the complex structure of knowledge generation and diffusion.

A final point, which however should not be seen as a limitation of the above-discussed model, refers to the fact that bilateral barter exchange is just one possible way of depicting knowledge flows; other options have been discussed in the literature (see Chapter 2) and are worth modelling. In what follows three further studies which attempt to address some of these issues are presented.

Modelling Knowledge Exchange as a Costs and Benefits Comparison

The model developed by Cassi and Zirulia (2008) attempts to address the first of the five limitations discussed above. Basically, the authors develop a model in the framework of the one proposed by Cowan and Jonard (2004), but explicitly consider costs and benefits associated to bilateral barter exchange. They assume that: ‘economic agents, while being embedded in social networks, interact with their social contacts only if it is convenient for them to do so. In other

words, the use of networks can be conceived as an economic choice based on cost–benefit comparisons’ (Cassi and Zirulia 2008, p. 78).

Moving from this observation, the authors develop a model in which a population of \bar{N} rational and self-interested agents, active in an exogenously given social structure, can choose between the following two learning options: individually, by improving their personal knowledge; or socially, by interacting and exchanging knowledge with other individuals in their social neighbourhood. Note that within each learning episode, the two mechanisms are mutually exclusive; hence an economic choice emerges as individuals have a limited endowment of time and resources to allocate either to one or to the other learning mechanism.

Each of the \bar{N} agents is located on a static and exogenously determined graph and is endowed with a vector of knowledge composed by K different categories of knowledge. As in Cowan and Jonard (2004), agents aim at maximizing the average level of their knowledge in different categories. They do so by acquiring new knowledge either by means of face-to-face interaction or by means of individual learning: in each period one agent is chosen randomly, and at random he/she selects one of his/her neighbours. These two connected agents can choose to barter knowledge or engage in individual learning, and will always choose the most convenient option. The decision rule adopted by Cassi and Zirulia assumes that agents can learn in only one category during each period. Moreover, agents are rational but myopic, in that they maximize only current period pay-offs. When mutually beneficial barter can occur, the Pareto superior outcome is selected (which holds for both the agents involved and the economy as whole), in which agents completely exhaust the knowledge trading opportunities.

The novelty of this approach rests in the assumption that agents operating in a social network balance costs and benefits associated with both knowledge exchange and individual learning, and make a choice between these two alternatives. Hence: ‘[t]he network constitutes an opportunity because it enables an alternative way of learning to individual one; at the same time, it constitutes a constraint because it selects the subset of agents out of the entire population, with whom an individual can interact’ (Cassi and Zirulia 2008, p. 81).¹² The authors claim that this modelling approach provides a better representation of some stylized facts emerging from empirical studies on know how trading. Specifically, they maintain that

their model captures the following three recurrent specificities: '[f]irst, exchange of knowledge can occur as a barter; second, effective exchange depends on benefit and cost comparison; third, knowledge exchange occurs between socially connected individuals' (Cassi and Zirulia 2008, p. 99).

Cassi and Zirulia allow for exogenously determined changes in the network structure tuning the rewiring probability p as described in Watts and Strogatz (1998) and implemented by Cowan and Jonard (2004). Hence, the authors confront a regular network with a small world network and a random network. Moreover, for each network architecture the authors consider different situations associated to a different opportunity cost of individual learning (they do so by varying a parameter associated to the easiness of individual learning).

The simulation experiment provides evidence partially contrasting with the results obtained by Cowan and Jonard (2004). The small world does not emerge as being unambiguously the most efficient network structure.¹³ In fact, Cassi and Zirulia observe that for a low level of opportunity cost, networks with the lowest average distance maximize efficiency of knowledge exchange; for an intermediate level of opportunity cost, networks with relatively low average distance and relatively high average cliquishness maximize efficiency of knowledge exchange; and for a high level of opportunity cost, networks with the highest average distance maximize efficient knowledge exchange. However, the small world structure does appear to be the most equal in terms of knowledge distribution, since the knowledge accumulated locally is diffused relatively more quickly among the other agents. Therefore, a trade-off emerges between efficiency and equity. From this finding it follows that the opportunity cost of using the network is a key variable determining its optimal structure in terms of aggregate performance (Cassi and Zirulia 2008, p. 99).

Modelling Knowledge Transfer as Localized Broadcasting

The two models just described saw knowledge flows occurring via a bilateral barter exchange. A distinct situation exists where agents engage in knowledge transfer which is not based on a *quid pro quo* scheme. In this case we refer to pure knowledge diffusion mechanisms and, specifically, to knowledge transfer models (see Chapter

2). Cowan and Jonard (2003) proposed a model in which a population of N agents innovates and accumulates knowledge. Innovation is exogenous and equally likely to be undertaken by any agent. Knowledge is defined as a scalar which evolves over time as the agent innovates and receives knowledge diffused by other agents. In fact, innovators diffuse their knowledge (by means of broadcasting) to their neighbouring agents, who receive and (partially) absorb the transferred knowledge.¹⁴ The population of agents is heterogeneous in two respects: ability to innovate and ability to absorb.

As in the two models described above, agents are located on a graph which can take the form of a regular network, a random network or a small world network (in accordance to the value assigned to the Watts–Strogatz rewiring probability p). Moreover, along with these three fixed networks (that is, agent interacts always with the same subset of n neighbours) the authors consider a random diffusion scheme where each agent broadcasts his/her knowledge to n agents chosen at random from the entire population. This modification leads the authors to approximate an epidemic model of diffusion of the type described in the section above. As put by Cowan (2004), comparing this model with the fixed architecture one provides an indication of the value of a fixed network as a diffusion vehicle.

Analysis of the results revealed a key factor to be the magnitude of absorptive capacity. When absorptive capacity is low, small worlds dominate in terms of long-run knowledge levels. However, as absorptive capacity increases, the random world tends to dominate more and more strongly.

Absorptive capacity also affects the comparison between fixed networks and random, epidemic diffusion. With low absorptive capacity, a fixed network performs better than random diffusion, independently from the value of p . On the contrary, if agents' absorptive capacity is high, a random diffusion pattern overperforms almost all fixed networks. Hence, the value of a fixed network comes from the repetition of interaction, which is most valuable when conveying knowledge effectively takes more than a single interaction (Cowan 2004). As put by the authors: '[t]he structure assists in overcoming the difficulties of absorption, and naturally, as those difficulties diminish [that is, as absorptive capacity increases], the value from the structure likewise diminishes' (Cowan and Jonard 2003, p. 529).

Modelling Knowledge Transfer as Face-to-Face Diffusion

Morone and Taylor (2004a) developed a model of knowledge transfer based on the assumption that agents meet in their social network and exchange knowledge by means of face-to-face informal interaction. This model represents a further attempt to address some of the limitations associated with the Cowan and Jonard (2004) model of knowledge exchange discussed above. Namely, the authors tackle the static–dynamic network issue and the complete information issue, introducing a network structure which changes as a consequence of interactions, based upon a more plausible mechanism for forming connections between distant agents, and in which individuals build internal models of the expected gain from interactions with their acquaintances.

A population of \bar{N} agents is located on a grid of cells (modelled as a wrapped grid so that there are no edge effects). Each agent is initially assigned a random position in the grid, and interacts with her/his closest neighbours. Not all the cells of the grid are occupied by agents, and those occupied are occupied by only one agent. The initial social network is created by connecting an agent with all other agents located within her/his ‘neighbourhood’: the local environment is defined as the region on the grid that includes those cells adjacent in the four cardinal directions and within the agent’s visible range (that is, von Neumann neighbourhood structure, see Figure 3.4).

The network evolves as agents get to know about the existence of other agents from their neighbours. Information about the existence of other individuals is transmitted as follows: if the acquaintance selected for interaction is connected to other individuals of whom the agent is not aware, then a connection is made from the agent to the acquaintance of the acquaintance. If there is more than one acquaintance of the acquaintance, then just one of them is chosen at random. Thus, an agent’s acquaintances list grows slowly over time, introducing a dynamic element to the network topology. Each acquaintance has at the outset an associated strength (ranging between one and zero) which is a measure of the strength of the relationship from the agent to her/his acquaintance. Each time an interaction results in knowledge gain, the strength of the relationship is incremented.

Agents have preferences for interactions with those acquaintances with whom they have strong relations.¹⁵ It follows that the selection




Figure 3.4 Von Neumann neighbourhood with visible range equal to 3

mechanism is not based on the assumption that agents have prior information about other agents' knowledge levels, but rather on an internal model of preference represented by the strength of relationships. The learning dynamic of the model is quite simple. For each period, each agent selects an acquaintance for interaction following the criterion described above. Once the contact is activated the contacted agent freely passes on some of his/her knowledge to the contacting agent. The exact amount of knowledge flowing from the sender to the recipient agent is calculated applying a 'gain function' which stems from a conjectural 'knowledge diffusion rule'. Such a rule is based on the assumption that agents with similar levels of knowledge are more likely to have gain interactions than agents with dissimilar levels of knowledge and therefore will be more likely to interact. More precisely, the knowledge interaction rule allows three different situations: (1) if the distance in knowledge between the two agents is very high there is low gain from interaction; (2) if the distance is intermediate, there is high gain from the interaction; (3) if the distance in knowledge is very low, there is again a low level of gain.

As argued by the authors, such a rule is theoretically supported by the concepts of 'cognitive distance' and 'cognitive proximity'

developed by Nooteboom (1999). Cognitive proximity is essential for understanding. However, ‘there must also be novelty, and hence sufficient cognitive distance, since otherwise the knowledge is redundant: nothing new is learned. If we specify effectiveness of communication as the (mathematical) production of communicability and novelty, learning is most effective at a distance which is neither too large nor too small’ (Nooteboom 1999, p. 140).

The results of the model are divided into short-term dynamics and a long-term stationary state. Long-term convergence emerges as a necessary outcome of the model. Nonetheless, different degrees of long-term convergences were generated by the occurrence or not of learning exclusion mechanisms. From a theoretical point of view such mechanisms can be triggered by the presence of a gain function with a very low absolute maximum which, in turn, generates an ‘ignorance trap’ for those agents initially endowed with a very low level of knowledge. The authors observed that increasing the initial minimum knowledge endowment provided to less knowledgeable agents and leaving unchanged the maximum level of knowledge possessed by most knowledgeable agents (hence, reducing the knowledge gap between agents) reduced the probability of agents falling into the ignorance trap.

Also, short-term dynamics showed different behaviours depending on the initial knowledge gap. In the case in which the knowledge gap is relatively large, the population divides into three groups: a first group of knowledgeable agents, a second group of catching-up agents, and a third group of marginalized agents. Compressing the initial knowledge gap the picture changes drastically: nearly all agents manage to catch up, reaching the highest possible level of knowledge in a relatively short time-frame.

As far as the network architecture was concerned, the authors always observed, over the short term, the emergence of small world properties. An important implication of these results is that small world networks can be associated with high variance in knowledge as well as low variance. The key variable to discriminate between convergence and divergence in knowledge diffusion is the initial level of knowledge variance: if initially agents are endowed with extremely different levels of knowledge, even though they live in a small world, they are not able to converge towards the highest level of knowledge. Some agents quickly reach the highest level of knowledge while others (in general, those initially endowed with low knowledge)

stay ignorant. On the other hand, if the initial knowledge gap is reasonably small (that is, the group is reasonably homogeneous) the process of knowledge diffusion quickly converges to the highest possible level of knowledge. From this finding the authors conclude that small world architecture can facilitate the equal diffusion of knowledge only if some barriers to communication are initially removed. If this is the case, the small world properties speed up the catching-up process; otherwise short-term divergence emerges.

The authors conclude that the main findings of their model of knowledge transfer relate to the generation of ignorance trap mechanisms which, in turn, produce a vicious circle for those people who are initially endowed with low levels of knowledge, and a virtuous circle for those endowed with high levels of education. Hence, they draw a relevant policy implication suggesting that giving access to a minimum level of knowledge to everybody will serve not just as a social stabilizer but also as a powerful tool to avoid the occurrence of undesired exclusion mechanisms.

SOME CONCLUDING REMARKS

In this chapter we presented a survey of knowledge diffusion studies identifying two basic classes of models. A first class assimilates knowledge diffusion to innovation diffusion. This approach starts off with pioneering studies on diffusion which investigated the patterns through which a new technology is adopted by a population of homogeneous agents and evolved into more sophisticated diffusion models which, making use of game theory, accounted for heterogeneous agents with heterogeneous beliefs. Such later studies, although building on social networks in order to avoid the simplifying hypothesis of complete mixing of social system members and allowing for different innovations occurring at the same time, suffered from a conceptual limitation which rests on the dichotomous definition of learning: agents can either learn of the existence of a new technology, and have the possibility of adopting it, or they can stay in their initial state of ignorance and not adopt it.

However, learning from peers involves more complex interactions and evolves by incremental improvements. In order to take on board these considerations a higher degree of complexity needs to be incorporated into knowledge diffusion models. This was achieved with a

new class of models which used mainly simulation techniques. The change of techniques has allowed the exploration of a wider range of phenomena through models that are more complex and at the same time remove formerly restrictive assumptions and design choices. Moving along the path set by this second class of models, in the next chapter we will present an agent-based simulation model which provides an original attempt to establish the complex relations linking knowledge sharing patterns, firms' partnering and the innovative capability of firms.

NOTES

1. From these assumptions it follows the representative agent hypothesis common to most Marshallian models.
2. As discussed in Chapter 2, to the view of knowledge as a public good it has been opposed a view of knowledge as a club good, that is, non-rival and relational (the diffusion of which is based on reiterated personal interactions).
3. This line of reasoning leads some scholars to focus on risk and uncertainty rather than on the process by which people become informed about something. However interesting, such models of technology diffusion will not be addressed in this book as they go beyond the scope of this study. We refer interested readers to Paul Geroski's paper (2000).
4. The label 'epidemic models' refers to the fact that such models were originally developed to study the transmission dynamics of communicable diseases.
5. The following analysis is based on Mahajan and Peterson (1985).
6. Note that if $N(t_0 = 0) = 0$, equation (3.2) can be rewritten as

$$\ln \left[\frac{1}{\left(1 - \frac{N(t)}{N}\right)} \right] = at.$$

7. These models depart from the seminal work of Zvi Griliches. The author's PhD dissertation, 'Hybrid corn: an exploration in the economics of technological change' (also published, in a different version, in *Econometrica*, 1957) is indeed the cornerstone of much of innovation diffusion literature.
8. For a complete description of the main limitations of basic epidemic models see Mahajan and Peterson (1985).
9. Note that this argumentation is tightly related to the discussion developed in Chapter 2 on the geographic dimension of knowledge flow patterns.
10. Note that this simplifying assumption requires a double coincidence of wants within the period, for knowledge exchange to occur. However, as discussed above and in Chapter 2, most empirical literature showed how exchange of knowledge does not occur simultaneously. Typically, agent I releases a piece of knowledge to agent J today because I expects to receive useful knowledge from J tomorrow (Cassi and Zirulia 2008, p. 88) or, as put by Morgan (2004), I expects to be reciprocated in kind tomorrow.

11. As observed by Robin Cowan: '[a] corollary of this is that as knowledge travels along a multi-agent chain, from i to j to k and so on, the knowledge degrades. Thus transmitting knowledge over a long chains is costly not only in terms of time, it is costly in terms of the diminution of the quantity of knowledge' (Cowan 2004, p. 12).
12. Building on this theoretical assumption of the model, the authors claim that their model, with respect to the existing literature, locates at an intermediate position along the dimension of exogenous vs endogenous networks. As put by the authors: '[o]n the one hand, we assume the existence of a (social) network that is exogenous and time-invariant, and, consequently, independent of agents' incentives to barter knowledge . . . On the other hand, agents' choices endogenously determine the actual network, which is given by the subset of links that are activated by agents. This actual network changes over time and clearly depends on agents' incentives to barter knowledge' (Cassi and Zirulia 2008, p. 81). However, we do not share this view as we believe that their network preserves all the characteristics of a static architecture since when agents choose an opportunity (whether to interact or not), it always refers to the same neighbours.
13. Note, however, that this finding was mitigated in Cowan and Jonard (2004) by the sensitivity analysis, which showed how the small world dominance over other network architectures was subject to the absorptive capacity parameter – that is, the optimal network structure becomes more and more random as absorptive capacity increases.
14. As in the previous models the effectiveness of knowledge flows depends upon the absorptive capacity parameter. However, in this context the parameter could be interpreted as an indicator of the relevance of tacit knowledge. As put by the authors: 'failure to absorb all available knowledge arises because codified, broadcast knowledge needs to be interpreted, and this interpretation intimately involves tacit knowledge that the receiving agent is unlikely to have completely' (Cowan and Jonard 2003, pp. 517–18).
15. Note that this model is not constrained to have symmetry of relationships between agents. In fact, more prestigious agents (with higher levels of knowledge) are the object of strong relationships with more peripheral agents (with lower levels of knowledge), which may be unreciprocated or reciprocated only weakly.

4. Knowledge diffusion and innovation: an agent-based approach

In the field of evolutionary economics, agent-based modelling is now recognized as one of the most promising new tools of investigation. As discussed in Chapter 3, an agent-based approach allows researchers to capture dynamics and complexity present in a model. This is exactly what is required for studying processes of innovation and knowledge sharing. The objective is to understand better the relations between micro-processes (the decisions and behaviours of economic actors) and the emergence of stylized facts common across much of industry (relating to innovation and informal relations among firms) in the model output. Recently, there has been a growing amount of research targeting this very area (for example Gilbert et al. 2001; Pajares et al. 2003) using agent-based methodologies. Gilbert et al. (2001) suggest that it has proved difficult to analyse innovation dynamics with the traditional analytical tools and suggest, as an alternative, the need for ‘an abstract simulation model that could constitute a dynamic theory of innovation networks’ (Gilbert et al. 2001).

Following this line of reasoning, in this chapter we present an agent-based simulation model which provides an original attempt to establish the complex relation linking knowledge-sharing patterns, firms’ partnering and the innovative capability of firms. However, before presenting the model we will briefly review some of the recent literature on knowledge, innovation and firms’ partnerships.

FIRMS' INTERACTIONS, KNOWLEDGE AND INNOVATION

In order to start reasoning about the link between knowledge and innovation we should first understand which kind of knowledge is needed to innovate, who possesses such knowledge and how it can be acquired by agents of innovation. The simultaneous ongoing processes of knowledge deepening and knowledge widening – which leads to a growing specialization of competences, as well as to a general expansion of the range of available technologies – calls for new learning efforts from firms. Innovative firms need specialized knowledge, as well as more types of knowledge which increasingly lie outside the firm itself. However, because of its tacit component, knowledge, and especially new knowledge, can be difficult to acquire in the market, so firms seek some form of collaboration with other firms and/or institutions that possess the required knowledge and, on a reciprocal basis, are keen on sharing it. This trend is reflected in the growing number of strategic research and development alliances which, starting from the early 1980s, has been striking (Hagedoorn 2002). Having a portfolio of alliances acts as insurance for a firm wishing not to be taken by surprise by new technology developments and increases the innovative capability of firms (Cowan 2004). Hence, firms act to create links through which to access disparate and specialized knowledge needed to innovate.

Such bilateral links, when considered all together, introduce the concept of relational networks. These are evolving networks which 'consist of relationships connecting actors . . . that are cooperating in order to acquire resources that they may not themselves possess' (Forsman and Solitander 2003, p. 5). The acquired resources (that is, mainly tacit knowledge) are then used to innovate. Consequently, these kinds of relational networks can be characterized as innovation networks – that is, networks which promote and favour innovative behaviours. Such networks of firms should be considered as a central locus for the creation of industrial novelty and consequently knowledge should not be considered as a freely available public good (as suggested by Arrow 1962), but referred to as a local (technology-specific), tacit (embodied in specific competencies of firms) and complex (based on a variety of scientific fields and technologies) asset (Pyka 2002). These characteristics of knowledge are, in a nutshell, the driving forces for the creation of innovation networks.

This approach brings us back to the ‘knowledge and geography’ debate discussed in Chapter 2. Classical studies by economists such as Alfred Marshall pointed out, nearly a century ago, the relevance of geographical proximity for general economic development and innovating activities. As Marshall (1920) observed, there may be geographical boundaries to knowledge flows. It is well established in the derivative fields of research that learning occurs more effectively when firms are located physically close to one another. Proximity is therefore related to learning-based comparative advantages which, in turn, result in a higher propensity to innovate. This idea gained momentum with the National System of Innovation (NSI) approach to economic development which contends that innovation predominantly involves collaboration and the exchange of tacit knowledge at the interfaces between organizations (FORFÁS 2004). National Systems of Innovations are seen as the joint outcome of three levels of analysis: (1) the firm level, in which companies are seen as repositories of knowledge embodied into their operational routines; (2) the meso-economic level of networks of linkages between private companies and other organizations, which enhance the firm’s opportunities of improving problem-solving capabilities; (3) the national level, composed by the set of rules, social relationships and political constraints into which microeconomic behaviours are embedded (Cimoli and Della Giusta 2000).

Hence, networks emerge, at the meso level, as a viable strategy to improve firms’ ability to innovate and handle increasingly complex problems. This is particularly relevant in regimes where technology is of major importance and the impossibility for a firm of mastering in-house the whole required knowledge can itself lead to the emergence of large informal networks via self-organization (Pyka 1997).

In what follows we will present an original model of knowledge diffusion and innovation. The modelled processes of knowledge acquisition and integration, as necessitated for innovation, result in the formation of networks of firms that have successfully innovated together. A further point which is central to the model is the idea that knowledge is structured. This idea has been developed in previous work (see Morone and Taylor 2004b) and highlights the fact that knowledge requires knowledge to be assimilated (Ancori et al. 2000). New knowledge can be acquired through interactive learning as firms partner to innovate. Conversely, we assume that the process of individual learning, such as that typified by the work of internal

R&D laboratories, occurs only at the initialization time of the simulation (this is a simplifying hypothesis which will allow us to concentrate only on the acquisition of additional knowledge provided by the external network).

A COMPLEX MODEL OF KNOWLEDGE DIFFUSION AND INNOVATION

Firms and their Social Network

The unit of analysis of this model is the firm. We assume a population of N firms allocated over a social network which is situated upon a grid of cells. Each firm is initially assigned a random position in the grid, and interacts with its closest neighbours. Not all the cells of the grid are occupied by firms, and those occupied contain only one firm. Initially, the neighbourhood is defined as the region on the grid that includes all cells adjacent and within the firm's visible range v . This arrangement is referred to as the Moore neighbourhood structure.

Figure 4.1 shows a Moore neighbourhood with $v = 3$. That is, firm A (placed in the central cell) is a neighbour of firm B because it is within three cells adjacent distance from B, where cells are permitted to be diagonally adjacent as well as horizontally and vertically adjacent.

The social network within which firms' interaction takes place could be represented as a graph where vertices correspond to firms and edges are firms' connections. Hence, we can write: $\mathcal{G}(I, \Gamma)$, where $I = \{1, \dots, N\}$ is the set of firms, and $\Gamma = \{\Gamma(i), i \in I\}$ gives the list of firms to which each firm is connected. This can also be written $\Gamma_x = \{y \in I \setminus \{x\} \mid d(x, y) \leq v\}$, where $d(x, y)$ is the distance from firm x to firm y , and v (visibility) is the number of cells in each direction which are considered to be within the firm's spectrum. Intuitively, Γ_x defines the neighbourhood of the firm (vertex) x .

The unit of time we define in our model is called the time step. In each time step, all firms are sorted into a random order, and then each is permitted to interact with neighbouring firms. However, this model introduces the possibility for each firm to acquire new neighbours (we will clarify the way in which this happens later in this chapter). This crucially implies that, as time passes, the social network evolves. Firms allocated in this social network aim at




Figure 4.1 Moore neighbourhood with $v = 3$

innovating. Innovation is defined as product innovation, that is, any time an innovation occurs a new product is supplied in the market. Innovation motivates firms to partner with other organizations (generating clusters of firms). In order to accomplish a new production process, new knowledge is required. When two (or more) firms partner together, let us say firm A and firm B, they integrate their knowledge (that is, they temporarily combine their specialized knowledge). Also, as described later, some knowledge transfer (from firm A to firm B and vice versa) takes place through interactive learning.

Defining the Firms' Skills Universe

The system is initially endowed with a Firms' Skills Universe (FSU), which contains the whole knowledge of the system.¹ In this model, the FSU is represented by a network of nodes and links: nodes in the FSU can be thought of as possible skills or technologies to be learnt by the firms, and links define the requirements of each node. The FSU structure therefore defines the way in which subsequent skills depend upon the prior acquisition of other skills. Using a similar




Figure 4.2 Firms' Skills Universe

graph notation to that used to describe the social network we can write: $FSU(\Sigma, \Psi)$, where $\Sigma = \{N_0, N_1, \dots, N_{MAX}\}$ is the set of skills, and $\Psi = \{\Psi(i), i \in \Sigma\}$ gives the list of requirements to go from one node to another.

In Figure 4.2 we reproduce a representative graph of a Firms' Skills Universe composed of 50 nodes (skills). This figure has been generated by employing a rather complex algorithm involving the definition of parent lists and child nodes. Through an iteration process, child nodes are subsequently added to parent lists that, with a certain probability, can split and therefore generate several independent areas of knowledge. We shall not describe the algorithm in detail here. It should be noted though, that since parent lists are allowed to split into independent parts, the resulting FSU will be composed of several branches which represent different fields of expertise in which individual firms can specialize.

Radical Innovations vs Incremental Innovations

The system is also endowed with an *ex ante* determined Global Innovation List (GIL), which represents all the possible innovations that can be achieved by firms. All of the potential innovations

are generated in the initialization phase of the simulation, at the same time as the FSU is created. At initialization, the number of innovations is specified, and they are then generated in sequence. This introduces some path-dependency in the construction of the GIL.

In the model, innovations take the form of vectors of skills. Building on the Schumpeterian tradition, two types of innovations are defined: incremental innovations and radical innovations. The first kind of innovation relates to the so-called ‘creative accumulation’ and is always based upon an already existing innovation, and is created by replacing one existing skill with a child, or a child of a child, and so on (we should say, closest available descendent) of that same skill. On the contrary, radical innovations, which relate to the ‘creative destruction’ process, are created by combining a whole new set of skills never used for previous innovations.

Definition of Firms’ SP

Each firm is initially assigned a Skill Profile (SP) which consists of a knowledge endowment – a subset of the universe of possible skills (the FSU). First of all, the firm is assigned the root skill, which is placed in its SP. The agent generates a target list of ‘child’ nodes of those already in its SP which are systematically acquired or learned. Recall that child nodes can be acquired only if all of its parents have been mastered (knowledge demands knowledge in order to be acquired). If the target node cannot be immediately learnt, then the algorithm backs up one level and first tries to acquire all of the parents of that node. Therefore, the assignment of firms’ SP takes place through a search process which goes from less to more specialized skills (that is, it is a depth-first search).

The new target nodes that are identified during the search process are added to the target list in a random order. Therefore, other than the depth-first assumption, there is no preference or goal for which nodes are acquired. As a result of the (non-goal-directed) search process that represents an initial learning or endowment phase, we generate a population of differentiated (that is, heterogeneous in terms of their SPs) as well as specialized (that is, they are eventually able to obtain some atoms of superior knowledge) firms. Note that this process takes place only at the initialization time of the simulation.

Main Simulation Phase

Having created the GIL, FSU and firms' SPs at initialization, the model is then ready to proceed to the main simulation phase where firms' partnering and innovation take place. The single objective of firms is to obtain all of the necessary skills to fulfil the requirements of an innovation. The first firm (or group of firms) to attain a particular innovation will be recorded as the first-mover (FM) innovator of that product innovation. In other words, that firm (group) is recognized as the first one to develop and market the product.

In this model, as we are only concerned with the innovation process, and not with the emergence of markets for those new products, when the simulation reaches the point where an innovation is accomplished and its FM identified, that innovation will be 'tagged' by the innovator in the GIL. Hence, an innovation is credited to only those firms performing it first.

Firms' Innovations and Partnerships

As already mentioned, the goal of each firm is to innovate. At every time step, agents are sorted into a random order. In its turn, a firm will select one element – an unaccomplished innovation – from the GIL. The firm will try to perform an individual innovation, by comparing the individual firm's SP (that is, the current possessed skills), with the selected vector. If successful, in other words if the firm possesses all of the required skills, it becomes the FM.

However, if the firm is not able to innovate individually, it will try to partner with its acquaintances and jointly innovate. In selecting an innovation, all firms follow a strategy which is relatively intelligent: the selection must be of one for which the firm possesses at least one skill; the selection is random, but with the chance of selection weighted according to the number of skills of that innovation already possessed. Partnering happens through direct interactions among neighbour firms. A partner is selected at random from within the neighbourhood of the firm. The firm and the partner now try to achieve the selected innovation by integrating their respective skills. If individually the partners each possess some of the required skills, but in partnership they do not possess all of them, then the search process will continue in the subsequent time step. In this sense, the strategy of selection is persistent: the firm initiating the partnership

will, in subsequent time steps, contact another of its neighbours, until either the partners can together perform the joint innovation, or there are no more neighbours to contact.²

Interactive Learning Process

In addition to the initial endowment of skills, that is, the creation of a firm's SPs, there is an interactive learning process whereby firms can obtain further skills as a product of their successful partnering for joint innovation. In this step, some skills are diffused from the innovator's profile to the partner firm's profile, and vice versa, thereby increasing the total level of knowledge of the system. This is repeated for each innovator-partner pair if there are more than two firms involved.

The potential number of skills diffused equals to the number of skills contributed to the innovation by the partner. For each firm in the innovator-partner pair, the skills to be acquired through interactive learning are based on the skills contributed by the other. These contributed skills are placed into a temporary target list, and acquired by the same procedure described above (using the dependency rule). Note that it may be the case that no new skills are learnt by a firm during the interactive learning step.

Joint Innovation and Network Change

The network is gradually changed by the addition of new links between distant firms, in addition to the existing links between neighbouring firms. The number of links to be added is determined by multiplying the rewiring parameter p by the total number of links present in the system at the beginning of the simulation. The new links are added gradually, at the same pace as innovations are attained (that is, when 10 per cent of innovations have been attained, 10 per cent of links will have been added). This represents the assumption that as firms successfully innovate they become more widely connected and hence have more opportunities to partner.

Given the above model specifications, we aim at investigating the emergence of innovation networks. Under growing values of p , we hypothesize, innovation clusters will emerge and the industry will innovate at a faster pace. We shall test whether this could be a general result of our model or if other critical conditions need to be met for such an outcome.

SIMULATION EXPERIMENT: EXPERIMENTAL SET-UP AND RESULTS

We used the JAVA platform with the RePast (Recursive Porus Agent Simulation Toolkit, North et al. 2006) libraries for implementing the model and JUNG libraries (Java Universal Network/Graph Framework 2007; O'Madadhain et al. 2005) for analysis of the networks' data. We carried out repeated simulation experiments (batches), to identify different trajectories of model behaviour. Over the batches we then took averages for all the relevant indices.

In this set of simulation experiments we used the following fixed parameters: number of time steps = 100, number of agents = 40, grid size = 20, number of radical innovations = 60, number of incremental innovations = 60, number of skills per innovation = 5, total number of skills = 200. The number of skills possessed by each agent is drawn from a uniform distribution with $\mu = 45$ and $\sigma = 10$. The visibility parameter v was set equal to 2. Moreover, the rewiring parameter p was varied as values drawn from the set: $\{0.1, 0.3, 0.5\}$. Hence, we have three different simulation specifications, each specification was iterated 100 times (using the same random number seeds); we shall now present the findings obtained with the simulation experiment.

First, we refer to average values of each simulation specification and compare the system performance in terms of both individual and joint innovations. As expected, the industrial environment performance is positively correlated with the rewiring parameter (see Figure 4.3). As p increases, the overall number of innovations introduced, in the long run, increases. However, the gap among the three simulations is quite narrow, varying from about 85 innovations achieved after 100 time steps when $p = 0.1$ to a maximum of about 92 innovations when $p = 0.5$. Hence, there is approximately an 8 per cent improvement in the system performance raising p from 0.1 to 0.5.

Furthermore, it is interesting to note that initially the rate at which firms innovate is almost unaffected by the rewiring parameter. In fact, the innovating performance in all three simulations is quite similar in the first 20 time steps, when the speed of innovation is rather high. Subsequently, the pace at which firms introduce new products slows down considerably and the three simulations start differentiating in accordance with the rewiring parameter. Note that the system performance in terms of joint innovation resembles the




Figure 4.3 Total and joint number of innovations, by various levels of p


total innovations performance. This reflects the fact that in the case of individual innovations the rewiring parameter has virtually no effect upon the speed of innovation.

In order to corroborate this finding, the correlation coefficients between the rewiring parameter and the joint and individual innovations were calculated.³ The correlation between JI and p was found to equal 0.20, suggesting the presence of a positive, although not very large, correlation between the two variables. As expected, the rewiring parameter resulted uncorrelated with individually performed innovations. Additionally, we calculated the correlation between p and the average size of innovating partnership (see Table 4.1) which resulted as marginally correlated. Interestingly, we found a much larger correlation between the number of joint innovations and the average number of agents involved in a partnership. In this case the coefficient is positive and equal to 0.66, suggesting that the larger the partnership is, the higher is the ability to jointly innovate. This is a crucial finding which suggests more thorough investigation of partnership formations.

We shall attempt to do so by analysing single runs in more detail.

Table 4.1 Correlation coefficients

	p	II	JI	<i>Avg. par.</i>
p		-0.006	0.203	0.125
II			0.040	0.049
JI				0.657

Figure 4.4 Innovation performances (rewiring parameter $p = 0.1$)



In order to underpin the key factors affecting innovating performances we will compare two runs which produce rather different innovating patterns. Specifically, we compare what we could label a ‘good run’ (run 90) with what we label a ‘bad run’ (run 98). What qualifies the run as good or bad is the total number of innovations performed after 100 time steps. Hence, a good run is one where a high number of innovations are performed and a bad run is one that is not very effective in terms of innovations.⁴

First, we compare the innovation performance in the two runs. We start by looking at the total number of innovations and the number of joint innovations achieved in each run, when the rewiring parameter is set equal to 0.1 (see Figure 4.4). As we can see, the good run overperforms the bad run mainly because of the higher ability to perform joint innovations. Hence, we should look at the network characteristics which facilitate the creation of collaborative partnerships.

In doing so, we start by looking at the density of the acquaintances network in the two runs.⁵ As clearly emerges from Figure 4.5 (top left panel), the good run is characterized by a much denser acquaintances network, resulting in higher opportunities for interactions. This finding is confirmed by the density dynamic of the partnership network which grows faster and converges to a higher level in the good run (Figure 4.5, top right panel).

If we look more thoroughly at the partnership network we can underpin some other differences between the two runs. For instance, when looking at the cliquishness of the two graphs (Figure 4.5, bottom left panel) we notice that it grows at a similar fashion in the first seven time steps; however, in the following nine time steps it drops in the good run and keeps growing in the bad run. After that, it grows slowly and steadily in both runs until it reaches the steady state equilibrium at values 0.39 and 0.45 respectively for the good run and the bad run. Hence, the cliquishness dynamic is non-monotonic in the good run. The slump in the pattern could be determined by the sharp increase in the average number of connections (depicted in Figure 4.5, bottom left panel) observable between time step 3 and time step 17. Moreover, in the long term the bad run reaches a superior equilibrium with respect to the good run, suggesting that the presence of a more clustered network (that is, one with higher cliquishness) does not necessarily reflect a more innovative environment. Above all, what really seems to affect the system performance, in terms of achieved innovations, is the density of the network and the average number of contacts upon which each agent can rely in order to initiate a new collaboration. The poor performance observed for the bad run is, therefore, mainly imputable to the fact that it is a more sparse network with, on average, a smaller number of connections per agent.

This finding is reflected in the size of the largest component.⁶ As shown in the bottom right panel of Figure 4.5, the size of the largest component (both for the acquaintances and the partnership networks) in the good run is almost three times bigger than that of the bad run. Basically, this finding suggests that the network structure in the good run offers a much wider set of opportunities for interactions than the one offered in the bad run. In fact, after about 40 time steps a member of the largest component in the good run is connected (that is, has a finite path length) with more than 30 agents and could join in a partnership composed by up to 20 agents.






Figure 4.5 Network statistics for good and bad run (rewiring parameter $p = 0.1$)

In the case of the bad run, the largest subset of connected agents equals ten and does not increase over time. Moreover, the largest component in the partnership network reaches seven at the most. This is implying that the opportunities for interactions (and, therefore, joint innovating) are much slimmer in the bad run.

We could conclude that it is really a scale effect that determines the diverging performances in the two runs. The scaling factor affecting the innovation patterns is the number of connections present in both networks. Such connections play a vital role in facilitating partnerships and joint innovation activities. The two networks, while being built using the same parameters, differ in some critical initial conditions: looking at Figure 4.5 we can immediately observe that the acquaintances network of the bad run, compared to that of the good run, is characterized by an initial lower density and a smaller largest component. Indeed these differences in the initial structure of the network shape the innovating patterns. Further insights will be provided by looking at the structure of the partnership networks which characterize the two runs examined. However, before stepping into the networks analysis we will briefly present the results obtained in the good and the bad runs varying the rewiring parameter.

First, looking at Figure 4.6 we can observe that the good run is always outperforming the bad run in terms of innovations. It is also noticeable that by increasing p , the gap between the bad and the good run grows wider. Interestingly, increasing the rewiring parameter exerts an opposite effect on the performance of the good and the bad run. In fact, in the former case, higher values of p are associated with a growing number of innovations achieved; whereas in the latter case, growing values of p are inversely related to the innovative performance.

These results are apparently surprising as one would foresee that the addition of new links between distant firms would benefit more the type of network emerging in the bad run – that is, one characterized by a low density and an initially small largest component. However, we should recall that new links are added gradually, at the same pace as innovations are attained (that is, when 10 per cent of innovations have been attained, 10 per cent of links are added). Therefore, the faster is the initial speed of innovation, the faster will be the network growth. On the other hand, a network which initially innovates at a rather slow pace could get ‘locked in’ to an





Figure 4.6 Innovation performances

underperforming pathway. Hence, recalling what we said above, increasing the rewiring parameter seems to improve the system performance only if the network is initially sufficiently dense and there is a sufficiently wide largest component. This result reflects the fact that in a dynamic network, composed of a sufficiently large number of successful innovators, firms become more connected and hence increase their opportunities to partner.

In other words, when increasing the value of p there are two possible outcomes. On the one hand, there is a positive feedback effect: a network which is initially sufficiently well connected will perform well in terms of innovations. This will quickly increase the number of connections among firms which, in turn, will enhance their innovative performance reinitiating the positive feedback process. On the other hand, there is a lock-in effect: a network which is initially highly disconnected will perform poorly in terms of innovations. This will undermine the available opportunities to increase the network density and, in turn, will impact negatively on the innovating performance. The two outcomes are exemplified by, respectively, the good run and the bad run. As shown in Figure 4.6, the good run simulation attains nearly all possible innovations in the GIL and, therefore, 30 per cent and 50 per cent of new links are eventually added (see Figure 4.7). Conversely, the bad run attains less than 50 per cent of innovations, and therefore far fewer additional links are added and density does not increase very much. The feedback between network density and innovation performance is strengthened with increasing p , resulting in a wider gap between good and bad run dynamics.

Figure 4.8 reports the growth pattern of the largest component. In this case also there are significant differences between the good run and the bad run. However, as opposed to the case in which the rewiring parameter was set equal to 0.1, with p equal to 0.3 and 0.5, the number of acquaintances shows a growing trend also in the bad run. Specifically, with $p = 0.3$ the number of acquaintances in the largest component grows from 10 to 11 and with $p = 0.5$ it reaches 18. Much higher values are obtained in the good run: now the number of acquaintances varies between 36 and 39 respectively for a value of p equal to 0.3 and 0.5.

We will now take a further step in our investigation, again using the JUNG libraries, to visualize the evolution of innovating networks. The following figures are some images which were displayed





Figure 4.7 Density of the acquaintances network graph





Figure 4.8 Number of vertices in the largest component (acquaintances network)

and instantaneously captured over the course of the simulation experiments, that is, they are snapshots.

The configuration of the network at different time steps over the course of the good run and the bad run simulations will be compared. In these simulation runs, the analysis will concentrate on the first 50 time steps (that is, before the system reaches convergence), looking at snapshots before and after critical periods in the acquaintance network evolution to explain better the innovation performance and development of the partnership network. Links present in the partnership relation are a subset of those present in the acquaintance network: $\mathcal{G}(I, \Gamma_{\text{partnership}}) \in \mathcal{G}(I, \Gamma)$. As explained earlier, the differences between good and bad run are due to the different initial location of the firms, whereas all other parameters are the same.

Figure 4.9 shows the acquaintance and the partnership networks (left panel and right panel, respectively) for the good run with $p = 0.1$ at time step 5. The left-hand panel shows a relatively dense network with a wide largest component (compare with Figure 4.5). These networks are hardly altered from the initial condition – the acquaintance network contains just one additional, distant, link between vertex 0 and vertex 16. The partnership network displays a number of partnerships already activated at time step 5. Moreover, the network architecture displayed in this figure does not vary significantly when changing the value of p .

Figure 4.10 shows the evolution of the acquaintance network graph in the good run for different values of p . Here, the networks are illustrated in a series of snapshots: time step 15 left-hand column, time step 30 centre column and time step 50 right-hand column.

Comparing the three rows one can see how the network becomes denser as the value of p increases. In fact, the rewiring allows links to be made between distant parts of the graph and offers more potential partners, as illustrated in Figure 4.11.

Figure 4.11 shows the evolution of the partnership network graph in the good run for different values of p at different time steps. Notice how the rewiring of the acquaintance network graph alters the opportunities to network and hence the shape of the partnership network. Interestingly, many of the links that are present in the partnership graphs for low values of p are not present at a higher value of p , where distant partnerships have become more important for innovation. For example the link between vertex 5 and vertex 34, the link between vertex 2 and vertex 15, and the link between




Figure 4.9 Acquaintance and partnership network graphs: good run snapshot ($p = 0.1$, time step 5)

vertex 18 and vertex 10, do not appear in the network for $p = 0.5$. This final network ($p = 0.5$, time step 50: bottom right-hand panel) appears to contain almost as many distant partnerships as it does local partnerships.

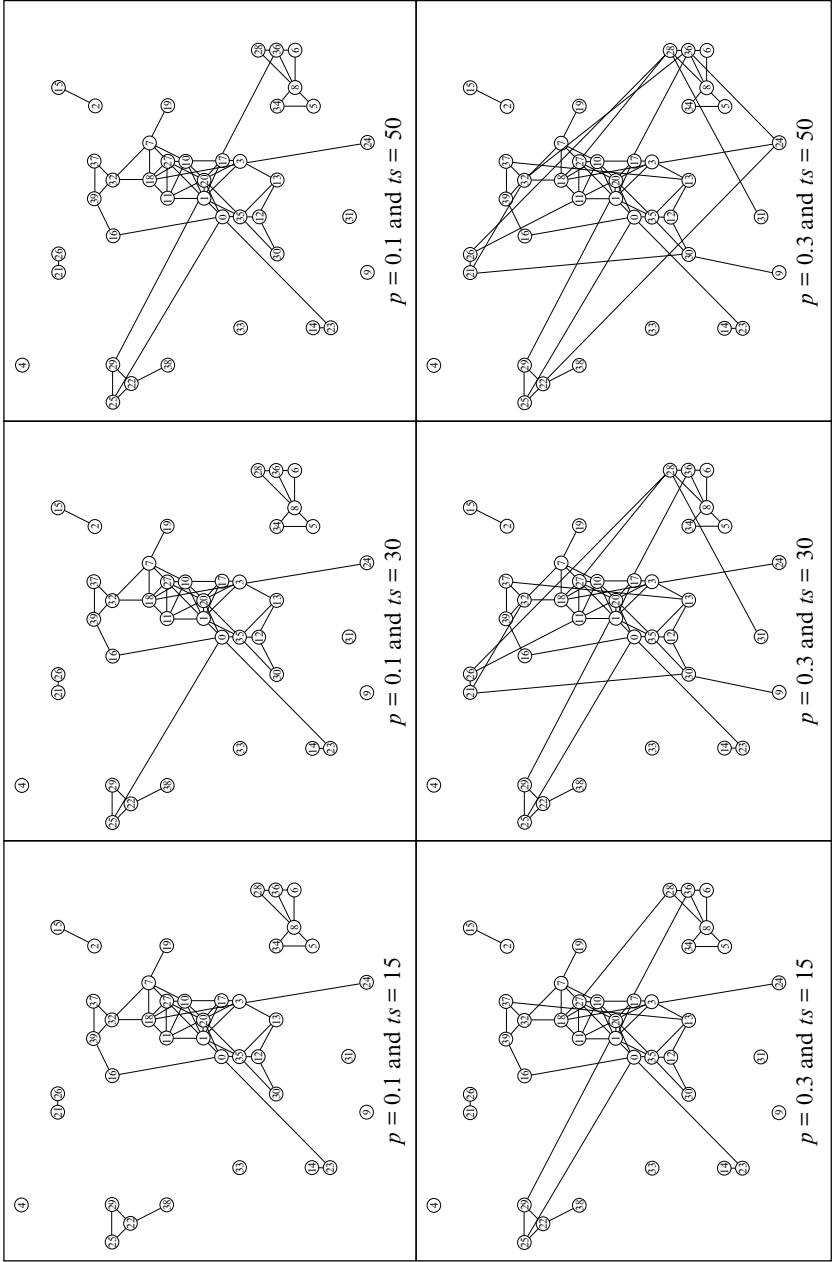
In Figure 4.12 we depict the acquaintance and the partnership networks (left panel and right panel, respectively) for the bad run with $p = 0.1$ at time step 5. As we mentioned in the corresponding good run network (Figure 4.9), at this stage there are no additional links added. Moreover, the network architecture (of both acquaintance and partnership graphs) at time step 5 does not vary significantly when changing the value of p .

Similarly to Figure 4.10, Figure 4.13 shows the evolution of the acquaintance network graph in the bad run for different values of p at different time steps. The first simulation performs very poorly in terms of additional links: there is only one link added between vertex 3 and vertex 37. In this case, innovation remains low, as shown in Figure 4.4. This picture improves marginally when increasing the value of p (bottom panels); though the number of available links (that is, the space for possible partnership) is much lower when compared to the good run. Indeed this is reflected in the partnership network displayed in Figure 4.14.

As we can see, the number of partnerships is rather low if compared with the good run performance displayed in Figure 4.11. This result is invariant for increasing values of p , confirming our earlier finding that increasing the rewiring parameter improves the system performance only if the network is initially sufficiently dense and there is a sufficiently wide largest component. Hence, the bad run appears to be locked in to an underperforming pathway.

Compared to the case of the good run, the partnership networks of the bad run are much more sparse and are not able to take advantage of additional, distant links in the acquaintance network. From the network snapshot illustrations, it can be concluded that firms in the good run are able to take advantage of opportunities with distant connections, whereas firms in the bad run are not able to do so. This is immediately visible in the network figures, where partnerships span all parts of the grid in the good run simulation with $p = 0.3$ and $p = 0.5$ (time step $t = 30$ and $t = 50$ in Figure 4.11) whereas in the bad run the graphs remain mainly unconnected (Figure 4.14).

In the good run, acquaintance networks are more effectively mobilized as a resource of the firm. The initial network architecture




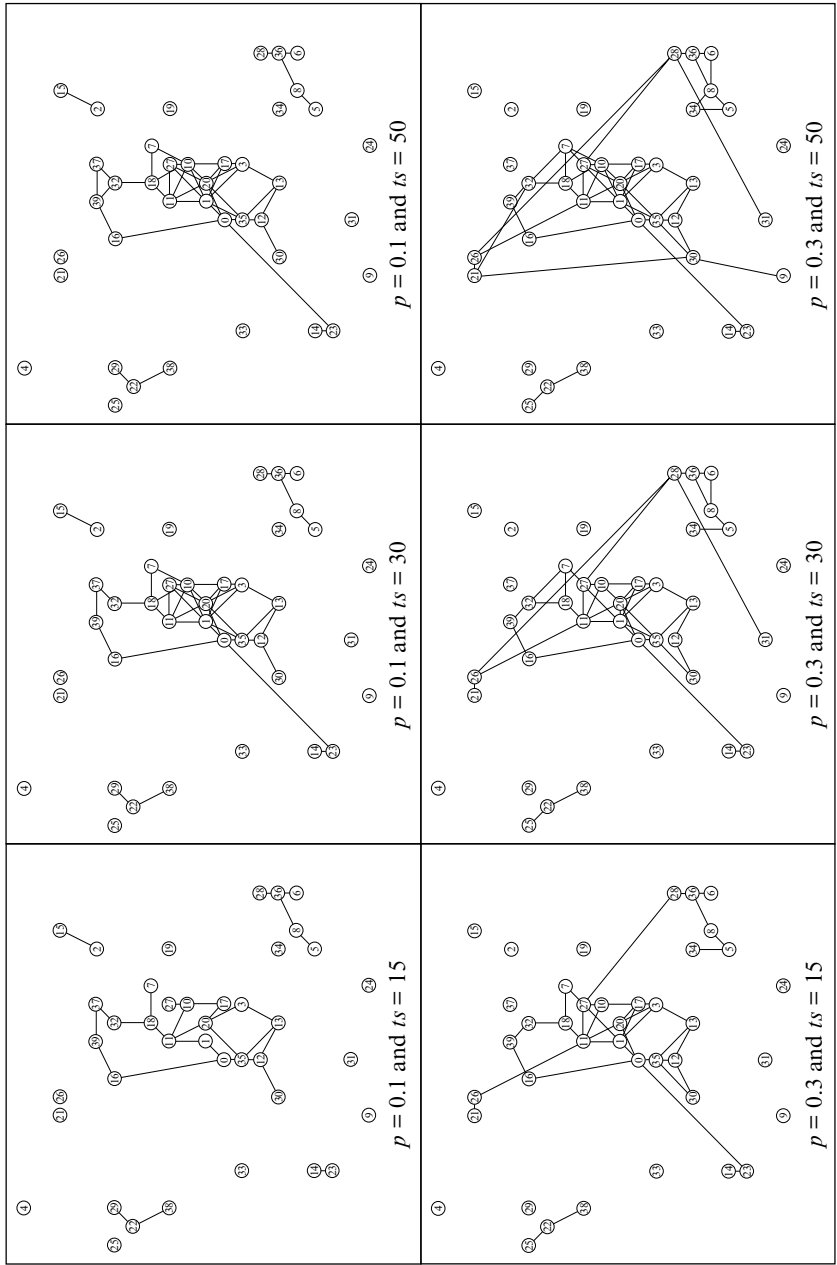


Figure 4.10 Acquaintance network graph: good run snapshot, various values of p and various time steps (ts)






Figure 4.11 Partnership network graph: good run snapshot, various values of p and various time steps (ts)

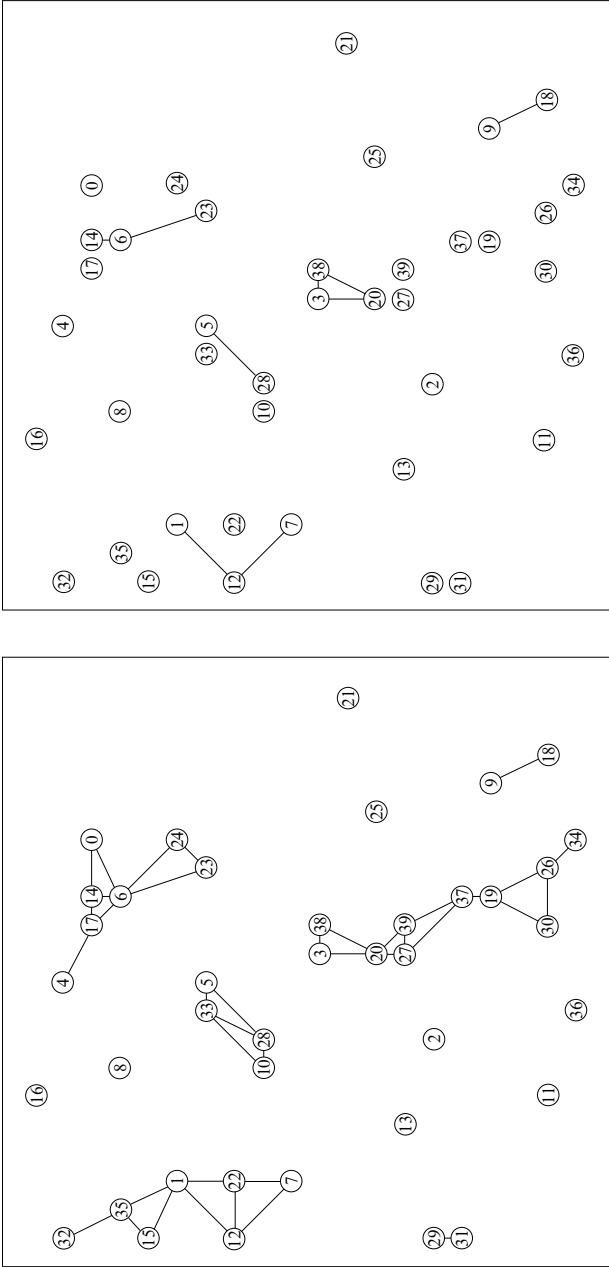


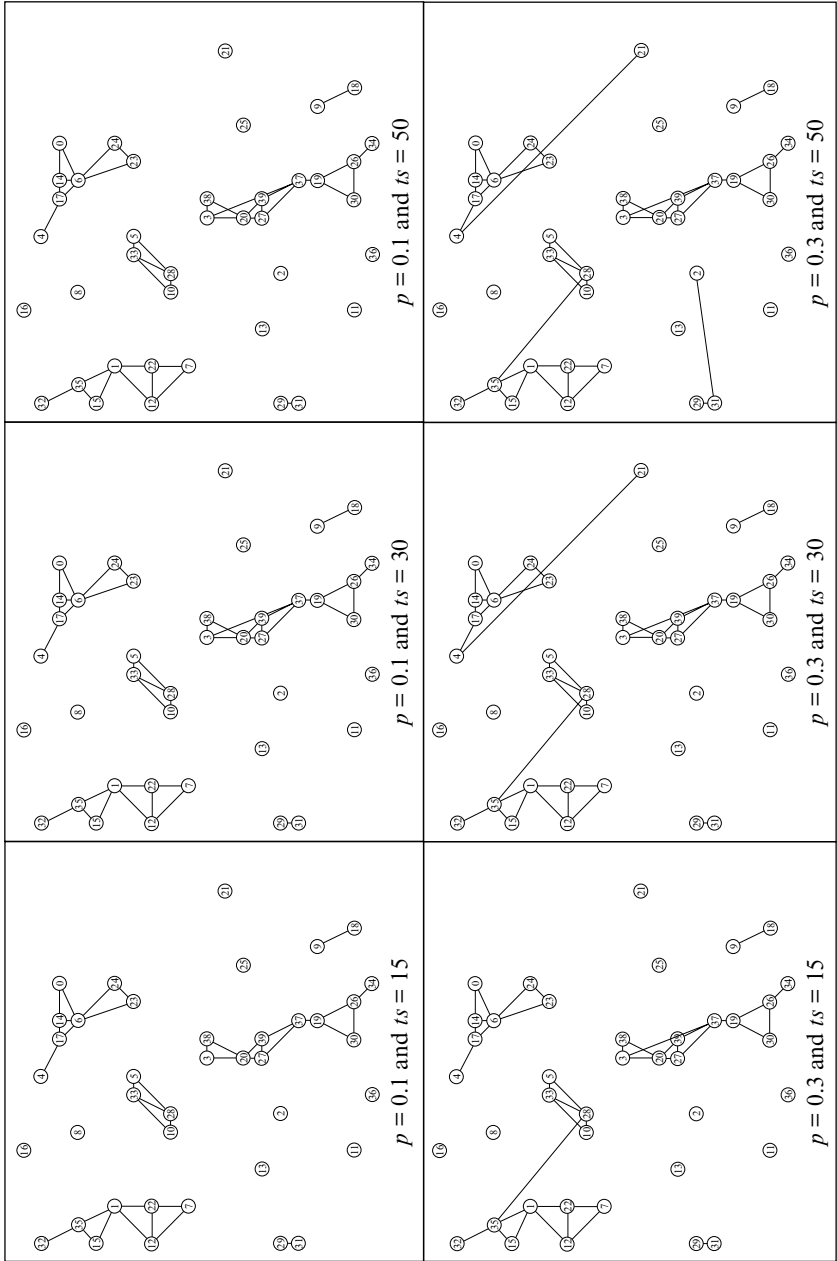
Figure 4.12 Acquaintance and partnership network graphs: bad run snapshot ($p = 0.1$, time step 5)

represents the space of opportunity for innovating and, as innovations are attained, this leads to more dense network arrangements, with distant connections between innovating firms playing a crucial role. It also results in interactive learning processes. Conversely, in the bad run, the initial conditions (in terms of both density and largest component size) seem to impede the network from evolving and taking full advantage of possible distant connections.

All in all, the network analysis here presented has confirmed earlier findings, suggesting that a high rewiring parameter represents just a chance for partnering and innovating with distant firms. A chance which under unfavourable initial conditions (that is, low density coupled with a relatively small largest component) fails to produce any positive effect.⁷

In order to assess the robustness of our findings we shall look at the whole set of simulations (that is, all 100 runs) and see if the results obtained looking at two single runs can be generalized. From the analysis conducted above we concluded that initial conditions, in terms of both density of the acquaintances network and size of largest component, affect the final outcome of the system – that is, the overall innovative performance. This finding was largely independent of the value of the rewiring parameter, hence suggesting that an initial low density and/or low largest component might trap the system into an underperforming equilibrium. We test this hypothesis by plotting the initial conditions against the steady state equilibrium. Specifically, we take the average performance (over the three simulation specifications) at the end of each run ($ts = 100$) and plot it against the initial ($ts = 1$) density of the acquaintance network and size of its largest component (note that both the density of the acquaintance network and the size of the largest component are invariant for different values of p at $ts = 1$).

Looking at Figure 4.15 we can observe that in both graphs there exists a positive correlation between initial conditions and final outcome (see the positive slope of the trend lines in both graphs), which suggests that those systems initially endowed with a low density and/or a small largest component are likely to be underperforming if compared to systems with more favourable initial conditions.




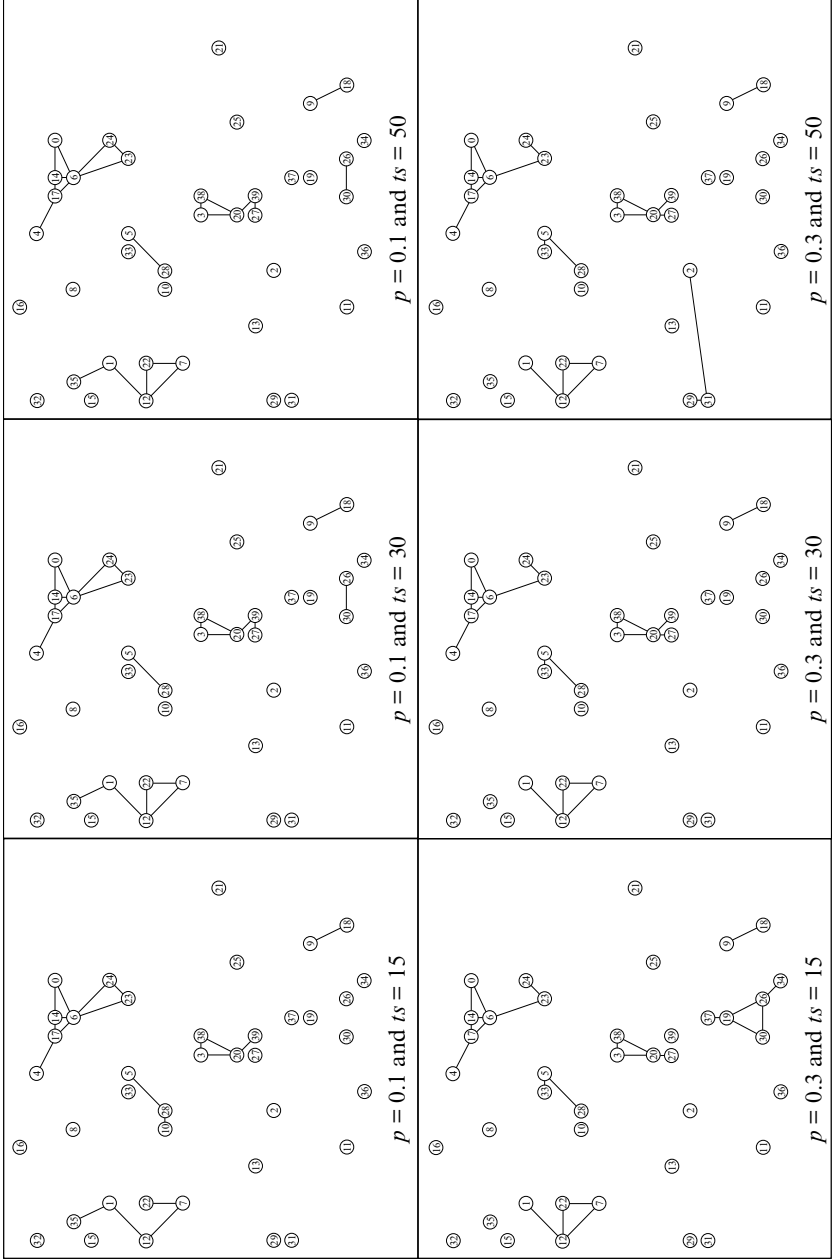


Figure 4.13 Acquaintance network graph: bad run snapshot, various values of p and various time steps (ts)






Figure 4.14 Partnership network graph: bad run snapshot, various values of p and various time steps (ts)





Figure 4.15 Impact of initial conditions over innovative performances

CONCLUSIONS

In this chapter a knowledge diffusion model with heterogeneous firms was presented. Firms were located randomly on a grid and were endowed with idiosyncratic competences (depicted as different skill profiles) which could be used individually (that is, by a single firm) or jointly (that is, integrating them in partnership composed by two or more firms) in order to perform innovation. Firms that innovated in partnership had the opportunity to acquire new skills (learn) by interacting with their partners. Moreover, we deliberately framed the simulation experiment in order to allow a feedback mechanism where successful systems (that is, those performing a sufficiently large number of innovations) became more widely connected. In this way, the structure could evolve into more dense network arrangements as innovations were achieved, simulating the emergence of innovation networks.

In this model the network has been specified as a knowledge resource, and knowledge integration the process by which the resource can be applied to innovation. This corresponds to the knowledge-based view of the firm, where firms are perceived as distributed knowledge systems which are required to integrate specialized knowledge efficiently, both internally and externally.

The results of the simulation exercises presented in this chapter have shown that the initial architecture of acquaintance networks is a most important factor for innovation and learning to occur. Although initial skill profile endowments are important, it is more than simply equating endowments of initial SP with firms' performance. In fact, the performance of the system showed high sensitivity to the arrangement of the firms' initial locations on the grid. Moreover, the way in which acquaintance networks were mobilized emerged as a key determinant of learning and innovation patterns.

All in all, our simulation model has shown how the network structure within which firms operate could represent both an opportunity and a constraint to innovating performances: above all, what really seems to affect the system performance, in terms of achieved innovations, is the density of the network and the average number of contacts upon which each agent can rely in order to initiate a new collaboration. A sparse network with, on average, a smaller number of connections per agent leads to poor performances in

terms of innovations achieved, whereas a dense network provides the right environment for partnership to occur and effective innovation networks to emerge.

NOTES

1. Note that in this model we assimilate the concept of knowledge to that of skill. Although rather similar, recall that some authors have distinguished between these two concepts. For instance, in Chapter 2 we discussed the distinction between tacit knowledge and skill suggested by Senker (1993).
2. Firms' innovation performance is recorded by means of an innovation 'score'. A record is kept of the number of times an agent has succeeded in an individual innovation or a joint innovation as FM or partner.
3. All correlation coefficients have been calculated by pooling together the three batches. Hence, correlation coefficients were calculated over a sample of 300 observations, where each observation refers to a single run.
4. Specifically, we selected the good and the bad runs in the following way: first, the average performance at the end of each run was calculated (that is, the average was calculated over the three simulation specifications); then, all 100 runs were ordered in a list according to their average level of performance; finally, the good run was randomly selected from the top 10 per cent of the distribution (that is, among the ten best runs) and the bad run was randomly selected from the bottom 10 per cent of the distribution (that is, among the ten worst runs).
5. Recall that the density of a graph is a measure of its cohesion: a dense graph is a graph in which the number of edges is close to the maximal number of edges. On the contrary, a graph with only a few edges is a sparse graph.
6. The nodes in a disconnected graph may be partitioned into two or more subsets in which there are no paths between nodes in different subsets. Each connected subgraph is called a component. The largest of such subgraphs will be the largest component.
7. Note that by looking at the networks configurations displayed above another conclusion can be drawn. In the case $p = 0.1$ a comparison can be made between the good run and the bad run: in both cases the trend is relatively steady innovation, but the partnership network shows decreasing cliquishness by the time the good run simulation reaches time step 15 and the partnership graph shows a reduction in cliquishness of around 25 per cent (see Figure 4.5, bottom-left panel).

In this simulation, a relatively dense largest component at time step 5 (see Figure 4.9, right panel) gains three new members by time step 15 (see Figure 4.11, top left panel). These are vertices 24 (connected via vertex 18), 23 (connected via vertex 0) and 14 (connected via vertex 23). There are no additional links made among nodes that are already members of the component. Recall that cliquishness is at a maximum when a vertex is linked to a single other vertex (where it has a value of 1) and at a minimum when it has no links, or if it has two or more links but none of the adjacent vertices are connected directly to each other (where it has a value of 0); this action is responsible for the above-mentioned reduction in cliquishness of the acquaintance network (observed in Figure 4.5, bottom-left panel). This effect is not seen in the bad run and does not appear to be a general behaviour of the good run class of simulations.

PART II

Empirical Studies and Model Validations

5. Empirical studies on knowledge flows

This chapter deals with empirical studies on learning patterns and knowledge flows. As discussed in the first part of this book, knowledge flow is a rather complex phenomenon which can be broken up into several processes. In Chapter 2 we distinguished between knowledge gain and knowledge diffusion processes; the former referring to formal and controlled flows and the latter referring to informal and largely uncontrolled flows. This classification is also useful when one has to define empirical measures of knowledge flows. In fact, as discussed earlier on, knowledge gains relate to flows of disembodied knowledge which could be measured by using, for instance, data on technologies and patents trade (Arora et al. 2002).¹ However, the applied researcher faces a much harder task when attempting to measure informal knowledge diffusion processes. In this case there are few available data in the official statistics;² hence, extra efforts need to be undertaken in order to gather sound information on the direction and the intensity of such flows.

This second part of the book, being about empirical research, will shed some light on the advances in the literature on measuring informal knowledge flows as well as on the relevance of model validation as a tool for testing models against the real world. Specifically, in the following sections of this chapter we will concentrate our attention on a recent body of empirical literature which has developed some new and well-crafted techniques to measure informal knowledge flows. Then, in Chapter 6 we will present a methodological discussion on theoretical and applied models upon which we will elaborate on the issue of agent-based model validation. Finally, in Chapter 7, we will apply validation methodology to the knowledge diffusion model presented in Chapter 4 of this book.

MEASURING KNOWLEDGE DIFFUSION: A CHALLENGE FOR APPLIED RESEARCH

Most of the recent empirical literature on informal knowledge diffusion focuses on the relevance of localized knowledge spillovers for innovation. As mentioned in Chapter 4, since the seminal contribution of Alfred Marshall (1920), scholars have reasoned on the relevance of geographical and relational proximity in facilitating the circulation of new ideas among firms, institutions and various other actors, promoting processes of incremental and collective innovation (Giuliani and Bell 2005). In fact, as discussed broadly in Chapter 2, the relevance of proximity for new ideas to spread rests on the tacit nature of certain knowledge, which requires face-to-face interactions in order to be transferred.

Departing from these considerations, several empirical researchers have focused their attention on clusters and networks as the locus of knowledge diffusion. However, recent studies have stigmatized 'the role of fuzzy social relationships and ill-defined spillover mechanisms as the basis of knowledge flows and learning processes within territory-bounded communities' (Giuliani and Bell 2005, p. 48), and have consequently proposed more structured mechanisms that shape these flows and processes (for example Dicken and Malmberg 2001; Malmberg and Maskell 2002; Amin and Cohendet 2004).

Following these guidelines, several authors have attempted to measure knowledge flows in a more direct and reliable way. A major effort has been made in order to produce an accurate definition of the type of social relations which leads firms to cooperate and share knowledge. Social network analysis has provided researchers with a powerful tool to achieve such an aim.³ Cantner and Graf (2006), for instance, applied network analysis to investigate knowledge flows in the network of innovators operating in the city of Jena in Germany. Meder (2008) further investigates the determinants of cooperation agreements in research and development (R&D) activities, focusing on both geographical and technological proximity as key factors which facilitate knowledge sharing.

Further insights have been provided by Giuliani and Bell (2005) who used network analysis to investigate the structure of knowledge flows, focusing on heterogeneous cognitive characteristics of wine producers operating in the Colchagua Valley in Chile. Morone et al. (2006) followed a similar approach, analysing the diffusion

patterns of various types of knowledge (distinguishing among technical knowledge, law system-related knowledge and market-related knowledge) in the cluster of organic producers operating in the province of Foggia in Italy.

A different approach was followed by Berends et al. (2006), who investigated knowledge sharing in industrial research. In their study the authors did not use social network tools: rather, they relied on an ethnographical analysis, based on passive participant observation. This allowed the researchers literally to trace all knowledge diffusion patterns which occurred within a firm.

In what follows we will review these recent studies, pinpointing the major elements of novelty introduced in each work and their contribution to providing sound and accurate measures of knowledge diffusion. Far from being exhaustive, this review aims solely at providing the readers with some examples of recent and innovative approaches to the investigation of empirical knowledge flows.

MEASURING KNOWLEDGE DIFFUSION WITHIN INNOVATION NETWORKS

As mentioned above, in a recent paper Cantner and Graf (2006) studied knowledge diffusion patterns in Jena (Germany). This empirical investigation uses the tools provided by social network analysis in order to analyse the evolution of the innovator network. The authors use data on patents⁴ that were applied for at the German patent office and were disclosed between 1995 and 2001.⁵ To include all patents that are relevant for Jena the authors collected all patents where at least one of the inventors named on the patent resided in Jena at the time of application. Following this procedure they gathered information on 334 distinct innovators and 1114 patent applications. Altogether 1827 inventors (977 of whom resided in Jena at the time of application) were involved in the development of such patents, which covered 29 out of 30 technological classes.⁶

These data were used to build two types of innovator networks. The first network links innovators by the kind of technological knowledge they have created – that is, the authors define a simple measure of technological overlap counting the fields of research each pair of firms have in common. The second type of innovator network relates to the notion of knowledge transfer through personal relationships

and is built in two alternative ways – that is, firms or research institutes are related if scientists know each other through working on joint projects (cooperation) or move from one organization to the other (scientist mobility). All in all, the authors define three types of innovator networks: (1) the technological overlap network; (2) the cooperation network; (3) the scientist mobility network.

The technological overlap network is constructed by defining first the two-mode sociomatrix with rows representing innovators and columns representing technological classes; subsequently, the authors construct the adjacency matrix computed as the product of the two-mode sociomatrix and its transpose. This network is interpreted as the potential for cooperation as it defines the degree of technological proximity among actors. Cantner and Graf maintain that being part of such a network is a necessary condition for cooperation as actors share a minimum of common knowledge which is needed for understanding each other (Cantner and Graf 2006, p. 467).

A visual inspection of the technological overlap network suggests that larger innovators form the centre of this network. This finding does not come as a surprise as it follows from the fact that large innovating firms are more likely to be involved in several fields of research. An exception to this finding is represented by Jenapharm, a large specialized firm operating in the pharmaceuticals sector, and therefore on the periphery. Looking at the evolution of this network over the two time periods considered shows that some firms move towards the periphery as they follow a strategy of higher specialization. A key public institution (that is, the University of Jena), on the other hand, moves towards the centre of the network as it increases the range of research fields in the second period.

Looking at basic descriptive statistics of the technological overlap network in the two periods the authors find increasing cohesion. This is interpreted as a stronger focus on core competencies, where the activities of the central actors become increasingly important for the whole network.

As already discussed, the existence of a link in the technological overlap network does not imply any real knowledge flow among two firms, nor that the two firms are actually related to one another. This problem is partially overcome by looking at the two other relational networks, which are constructed following a procedure similar to the one employed in the construction of the technological overlap

network. By creating a two-mode sociomatrix where the innovators are the nodes (rows) of the network and the inventors on the patent are the characteristics (columns) of these innovators, Cantner and Graf identified those inventors that have worked on research projects for more than one innovator, thereby creating linkages among these innovators. As mentioned above, the authors consider two different criteria for a relationship to be established. The first criterion relies on direct cooperation: whenever the authors find a patent with more than one innovator (co-application), they assume it to be a cooperation. Their second criterion is less direct: '[i]f an inventor is mentioned on patents applied for by different, not co-applying innovators within one of the two periods of observation (1995–1997 and 1999–2001) . . . [we establish] a link between those innovators that is referred to as scientist mobility' (Cantner and Graf 2006, p. 470).

Looking at the evolution of these two networks it emerges that the knowledge among large, core actors increasingly flows through formal cooperation while smaller, surrounding or peripheral actors rely more substantially on informal, personal relations (captured by scientist mobility) as a means of knowledge diffusion.

Comparing the data collected over the two time spans allows one also to characterize the innovators according to their innovator status, that is, entry, exit and permanent. Studying the entry–exit dynamic for the network of technological overlap leads the authors to conclude that: 'the dynamics of the system tends towards an increasing focus on core competencies of the local innovation system; that is, innovators on the periphery of the network exit and new entrants position themselves closer to the core of the network' (Cantner and Graf 2006, p. 478). Finally, employing the network regression methodology the authors do not register persistent linkages through cooperation. This last finding is interpreted as evidence of the fact that actors do not tend to cooperate with previous partners. However, enduring relations and trust turn out to be relevant elements for scientist mobility: as workers or scientists change their jobs, they carry knowledge about competencies and the trustworthiness of previous colleagues or their prior employer (Cantner and Graf 2006, p. 478).

All in all, this paper adds to the empirical literature on knowledge flows as it uncovers some relevant mechanisms of cooperation and cluster formation (such clusters being the loci of informal

knowledge diffusion). Moreover, it sheds light on the relevance of scientists' cooperation and their mobility (which are indirect sources of knowledge flows) for innovation to occur.

Further insights on knowledge diffusion processes which occur through cooperation among innovating firms are provided by Meder (2008). The author presents an empirical analysis on the impacts of technological and geographical proximity on cooperative innovation activities, analysing the interplay of both dimensions. In particular, Meder attempts to answer the question of whether technological and geographical proximity affects the choice of the cooperation partner. As discussed in Chapter 2, both these dimensions of proximity are relevant for knowledge flows to occur. 'The degree of knowledge exchange and the success propensity of an R&D cooperation depend on the technological proximity between the potential cooperation partners . . . An increasing technological proximity facilitates knowledge exchange, which is one core incentive to engage in an R&D cooperation' (Meder 2008, p. 5). Moreover, Meder observes that the beneficial effects of geographical proximity 'seem to be due, in particular, to the possibilities offered by face-to-face contacts . . . [which] are required for the exchange of tacit knowledge which is, again, a core incentive to engage in R&D cooperation' (Meder 2008, pp. 5–6).

Similarly to Cantner and Graf (2006), Meder uses patents data drawn from the 'Deutsche Patentblatt' publication⁷ and examines the initializing conditions for engaging in interactive learning processes. Such conditions are examined by analysing the impact of different actors' characteristics on cooperative agreements in the field of R&D. Cooperation agreements refer to the year 2003 whereas independent variables refer to the period 1998–2002. This allows the author to introduce a dynamic element into his analysis as the willingness to engage in a specific R&D agreement today depends on factors developed in the past.

Meder performs a regression analysis using a dummy as the dependent variable which takes the value of one if the firm has engaged in a cooperative agreement in 2003, and zero if otherwise. In that year 1333 German actors filed for 1089 collaborative patents. This is a rather small proportion of the potential collaborative agreements which include all possible pairs of cooperation. The dataset ends up with 887 778 observations (possible pairs of cooperation), with 1089 observed cooperations.

The set of independent variables includes: technological proximity (defined as the technological differences among firms and constructed using information on technological fields reported on patents); geographical proximity (defined as the distance in space between two firms and measured in kilometres according to the postal code marked on the patent application); attractiveness of being cooperation partner (this variable accounts for the valuable knowledge which is offered by actors of a certain pair in the dataset and is calculated using information on the number of patent applications of the three years 2000–2002); former cooperation experiences (this variable uses information about the cooperation activities of actors in the five years between 1998 and 2003, distinguishing between pairs in which both actors have cooperation experience and pairs in which only one partner has previous experiences); public research agreements (defined as a dummy which has a value of one if at least one actor in a pair is identified as a public research actor).

Running a ‘prior correction’ logit model⁸ the author obtains some interesting results. First, the author observes that technological proximity and geographical proximity enhance (independently) the probability of being involved in a collaborative R&D project. Moreover, combining both dimensions of proximity does not exert an extra effect on the cooperation probability. Meder found also that pairs of actors where one actor had experiences in cooperation (at least one co-application for 2003) display a lower cooperation probability than pairs where neither actor had such experiences; and pairs where both actors had cooperation experiences in terms of co-applications in the past were more likely to file for a patent together than pairs without such experiences. Finally, it emerged that pairs of actors identified as public research actors have no higher probability to engage in a common R&D project than other pairs of actors.

The element of novelty of this paper (similarly to the Cantner and Graf study) rests on the use made of patent data in order to define cooperation. Meder builds a database which allows him to identify the influences of technological and geographical proximity on the probability of cooperation agreements in the field of R&D, such cooperation agreements being perceived as an indirect measure of knowledge diffusion. However innovative, Cantner and Graf (2006) and Meder (2008) both use somewhat indirect measures of knowledge sharing which inevitably casts doubts on the effective measurement of knowledge flows among cooperating agents.

More direct measures of knowledge diffusion are needed in order to gain confidence on the actual dimension of informal learning. In what follows we will present three studies which addressed this issue, using case studies largely based on fieldwork.

MEASURING KNOWLEDGE DIFFUSION WITHIN SOCIAL NETWORKS

By means of a fieldwork study, Giuliani and Bell (2005) investigated differences in firms' knowledge endowments and analysed how such differences influence the formation of intra- and extra-cluster knowledge networks. Among other things, the authors attempted to understand whether: (1) firms with higher absorptive capacity⁹ are more likely to establish knowledge linkages with extra-cluster sources of knowledge; and (2) such firms gain leading positions within the cluster. We shall concentrate our attention on these two research hypotheses. The relevance of extra-cluster knowledge relations rests on the fact that sole reliance on localized knowledge can result in the 'entropic death' of the cluster that remains locked in to an increasingly obsolete technological trajectory (Camagni 1991; Grabher 1993; Becattini and Rullani 1993; Guerrieri et al. 2001; Cantwell and Iammarino 2003).

As already mentioned, the empirical analysis was largely based on social network methodology and used primary data at firm level gathered via interviews (based on a structured questionnaire) in a sample of firms operating in the wine cluster. Interviews were directed to technical employees (chief oenologist or the cellarman), which provided reliable information about the history and current characteristics of the firms and proved to be key nodes in the cognitive interconnections between firms. Along with general background and contextual information, the data gathered in the interviews were used to develop a set of quantitative indicators in three key areas: (1) the 'absorptive capacity' of the firms; (2) their intra-cluster knowledge communication patterns; and (3) their acquisition of knowledge from extra-cluster sources.

Specifically, absorptive capacity was broken into four components (the level of education of the technical personnel employed in the firm; each professional's months of experience in the industry; the number of firms in which each professional had been previously


employed; and the type and intensity of R&D undertaken by the firm) which were summarized in an index employing principal component analysis. Intra-cluster knowledge communication patterns were measured using different centrality indices drawn from graph theoretical methods. Such indices include: (1) out-degree centrality index;¹⁰ (2) in-degree centrality index;¹¹ (3) betweenness;¹² (4) in-degree–out-degree centrality index.¹³ Finally, a measure of external openness (which allowed measurement of the acquisition of knowledge from external sources) was computed as the number of linkages with extra-cluster sources of knowledge.

Applying these techniques allowed the authors to define and measure intra-cluster knowledge flows and establish the relevance of extra-cluster knowledge sources. This, in turn, led the authors to assess the knowledge structure of the cluster and to provide answers to their research hypothesis – that is, whether firms with higher absorptive capacity were more likely to acquire extra-cluster knowledge, and whether such firms had a central position within the cluster.

Overall the network studied by Giuliani and Bell displayed the characteristics of an open knowledge system ‘as many of its constituent firms have established linkages with external sources of knowledge’ (Giuliani and Bell 2005, p. 54) such as leading research and technology transfer institutions and universities. Additionally, knowledge flows into the cluster from international sources; in particular, some local wine producers (labelled by the authors as ‘interface actors’ or ‘nodes of connections’) were engaged in knowledge relations with foreign consultant oenologists that play a major role in the transfer of frontier knowledge and techniques in the field (Giuliani and Bell 2005, p. 54). Along with what was postulated by the authors, the level of interaction with external sources of knowledge varied according to the absorptive capacity of individual firms.

Figure 5.1 (reproduced from Giuliani and Bell 2005) depicts intra-cluster knowledge diffusion patterns. Similarly to the network graphs presented in Chapter 4, this figure is quite informative as it shows the direction of knowledge flows (that is, the arrow of each tie) and the absorptive capacity of each firm (that is, the diameter of the node).

First and foremost, the authors observe that: ‘firms tend to interconnect differently to one another: in particular one group of firms (centre of the figure) are linked, transferring and receiving knowledge



Note: An arrow from I to J indicates that I transfers knowledge to J. The diameter of the nodes is proportional to firms' absorptive capacity.
 Source: Giuliani and Bell (2005).

Figure 5.1 Knowledge diffusion patterns in the Colchagua Valley

from each other. In contrast, another group of firms (top left) remain cognitively isolated' (Giuliani and Bell 2005, p. 57). Once they had established the existence of heterogeneous cognitive positions within the cluster, the authors attempted to test whether firms that were more cognitively interconnected in the cluster knowledge system also had higher absorptive capacities.

By means of a set of correlation tests Giuliani and Bell were able to underline the existence of a statistically significant relationship between firms' absorptive capacity and the different centrality indices discussed above. Their empirical findings showed clearly that average absorptive capacity varies considerably across the different cognitive positions. Particularly interesting results emerged from the correlation between the absorptive capacity and the in-degree–out-degree centrality index. As pointed out by the authors:

[t]his result supports the idea that a threshold for inter firm knowledge exchange exists, so that when firms' absorptive capacity is very low, the cognitive distance with other firms' knowledge bases becomes too high (i.e. infinite) and the firms tend to be isolated. Correspondingly, those firms that are sufficiently above the minimum threshold have a higher probability of being interconnected with other local firms. (Giuliani and Bell 2005, p. 58)

Combining the data about the external openness of firms and the cognitive position of firms within the local knowledge system, the authors were able to identify five main learning patterns within the cluster, which were summarized in five types of 'cognitive role': (1) technological gatekeepers (TG);¹⁴ (2) active mutual exchangers (AME);¹⁵ (3) weak mutual exchangers (WME);¹⁶ (4) external stars (ES);¹⁷ (5) isolated firms (IF).¹⁸ In conclusion, the study conducted by Giuliani and Bell makes a valuable contribution to empirical knowledge flows literature as it sheds light on the processes which control intra-cluster and extra-cluster knowledge diffusion processes. The methodology employed in their study, largely based on network and graph theory, has proved useful to associate heterogeneous firms' characteristics (mainly referring to the knowledge base of the firm) to the different role that individual firms play within the cluster. All in all, this study 'suggests that a cluster is a complex economic and cognitive space where firms establish knowledge linkages not simply because of their spatial proximity but in ways that are shaped by their own particular knowledge bases' (Giuliani and Bell 2005, p. 64).

As already mentioned, a similar study has been conducted by Morone et al. (2006). The authors studied knowledge diffusion patterns in a backward area located in the south of Italy (that is, the province of Foggia) and investigated the occurrence of knowledge flows in two different networks: (1) the network of local organic industrial firms; and (2) the network of firms and local institutions. Moreover, the authors argue that it is important to include institutions in such a network as they play a vital role in gathering information and knowledge from different sources (which, sometimes, lie outside the local network): some institutions produce knowledge by themselves (for example research institutes and research departments); others exchange information and knowledge among themselves and subsequently diffuse it to firms; and others diffuse codified knowledge obtained from legal and technical sources.

An initial set of organic producers was selected with the focus group technique. Subsequently, the firms' network structure was augmented following a free recall approach. The sample of local institutions included all institutions which supported organic food production. All in all, the authors questioned a sample of 66 firms and 16 institutions.

The questionnaire, submitted with face-to-face meetings both to firms and institutions, was structured in two parts. The first part aimed at gathering general information on the characteristics of the firm or institution. The second part aimed to collect information on relations and, more precisely, on the existence of ties and their nature.

First the authors defined the network of interactions, which contains all ties amongst the firms regardless of their nature (that is, whether they are a trade relation, information exchange or a longer-lasting cooperative relation). Subsequently the authors used the same actors to build a new network (which is a subset of the whole network of interactions) called the 'communicative network'. It contained only those ties identified as 'communicative exchanges of knowledge'. The communicative network was further broken down into three subtypes: technical knowledge that can affect directly the firm's productivity, organic production laws, and knowledge or information about markets and consumers.

The quite high level of desegregation of the analysis performed allowed the authors to investigate knowledge flows distinguishing among different types of knowledge. First and foremost, despite a

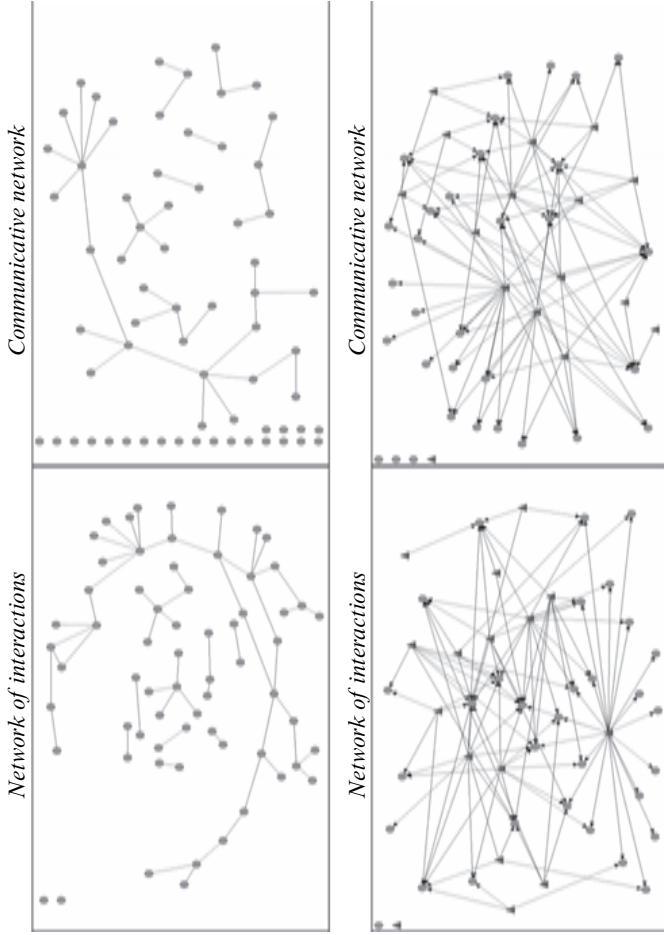
cohesive 'network of interactions' among organic producers, the authors observed that knowledge-based exchanges among firms were fairly marginal. Foggia's organic producers appeared not to be taking full advantage of social networks, perhaps due to negative attitudes about cooperation.

However, the authors found that such a communication void was partially filled by active local institutions (see Figure 5.2). In fact, the firms' institutions networks were much more dense and effective in diffusing knowledge. Specifically, the authors showed that institutions were more effective in diffusing juridical and technical knowledge (sector specialists such as agronomists, colleagues and organic farmers provide most technical support) and less effective in providing commercial or market information.

All in all, this study contributes to empirical literature on knowledge diffusion in various ways. First, it underlines the fact that proximity and the existence of social ties are just necessary (and not sufficient) conditions for knowledge to diffuse informally. In order for knowledge to spread effectively across a social network, the actors involved should be willing to engage in informal diffusion, and a certain degree of cooperation and trust needs to be present among actors. Moreover, this study shows that institutions can play a vital role in promoting knowledge flows. Finally, distinguishing among various types of knowledge allowed the authors to reach a finer grain of inspection and to point out which types of knowledge circulate more easily than others.

MEASURING KNOWLEDGE DIFFUSION IN INDUSTRIAL RESEARCH

A rather different approach was followed by Berends et al. (2006) who performed an in-depth analysis of knowledge diffusion patterns¹⁹ by means of two field studies conducted on two different industrial research groups.²⁰ Following an ethnographic approach, the authors studied the actual practices of researchers by implementing a passive participant observation: 'one of the authors temporarily shared a room with different researchers, followed them to meetings and to their laboratories, joined them for coffee and lunch breaks and on other social occasions, but did not actively participate in their research' (Berends et al. 2006, p. 87). Following this methodological



Note: Dots denote firms, triangles institutions.

Source: Morone et al. (2006).

Figure 5.2 *Network firms-firms and network institutions-firms*

approach allowed the authors to identify 227 episodes of knowledge sharing²¹ which were classified by means of a coding technique. This coding process eventually yielded three dimensions that were used to distinguish between mechanisms for the origination of knowledge sharing. The three identified dimensions of knowledge sharing refer respectively to: (1) the shared content (what is shared?): the authors distinguish between new content or existing content; (2) the source determining the sharing process (who is the actor who determined the content of knowledge sharing?): the authors distinguish among the person who is sharing his/her knowledge, the person that this person is sharing his/her knowledge with and the management which might determine the content of knowledge sharing; (3) the orientation of knowledge sharing (with what objective in mind is existing information selected or new information developed?): the authors distinguish among four possible orientations – orientation towards one's own problem (the sharing person's problem), orientation towards the other's problem, orientation towards a shared problem, or not oriented towards a particular problem.

These three dimensions can be combined in 24 logically possible knowledge-sharing mechanisms, out of which 16 were observed in the 227 episodes of knowledge sharing identified in the case study. Some of these mechanisms also involve, along with knowledge sharing, the creation of brand new ideas. Seven mechanisms were described in detail in the paper, also making good use of real examples drawn from the fieldwork which help in clarifying the occurrence of the actual knowledge-sharing processes.

The first mechanism is labelled 'diffusion' and it occurs whenever a member of an organization selects and communicates existing information without being oriented towards a particular problem. In this case the knowledge sharing is not meant to help anyone in particular. 'Information retrieval' is the second mechanism, which takes place when someone who needs a particular piece of knowledge or information obtains it by asking someone who has it. The content of knowledge sharing is thus not determined by the sharing person but by the other, and is oriented towards the other's problem. 'Information pooling' arises when the person sharing information chooses to do so because of a problem shared with others. 'Collaborative problem-solving' is a mechanism which consists of developing new information with regard to a shared problem (this is a case in which knowledge sharing is accompanied by knowledge creation).

The four knowledge-sharing mechanisms described above were found in the episodes observed in both field studies and correspond to models of knowledge sharing that are assumed in particular streams in the literature. Three other mechanisms that were frequently observed have received little attention in the mainstream literature on knowledge sharing and, therefore, represent an interesting element of novelty in the study of Berends et al. Such mechanisms have been labelled 'pushing', 'thinking along' and 'self-suggestion'. In the first of these originating mechanisms, pushing, the sharing person chooses to provide someone else with existing information. This mechanism is typical of gatekeepers, who monitor (external) developments and pass on to their colleagues what they think might be useful to them. Thinking along occurs when someone develops new ideas, hypotheses or questions with regard to someone else's problem, and it is not confined to informal meetings between two researchers. Finally, self-suggestion occurs when someone thinks about his/her own problem during interaction. The need to explain one's own problem or the need to defend one's own ideas stimulates a person to come up with new explanations, solutions, arguments and conclusions in the same way as people can think about someone else's problem. Note that both thinking along and self-suggestion involve new knowledge creation while sharing knowledge.

Along with the definition of the knowledge-sharing origination taxonomy, the authors provided an explanation as to why each mechanism is valuable for R&D and a detailed description of the preconditions required by each mechanism to be effective. Indeed, the depth of the study allows the authors to underpin several rather specific characteristics of knowledge diffusion. This is the major element of novelty introduced by this study; however, as acknowledged by the authors themselves: 'the concepts and findings discussed . . . should be tested and elaborated in further research. It should be explored whether the same origination mechanisms can be found in other organizational functions, such as engineering and marketing' (Berends et al. 2006, p. 94).

CONCLUSIONS

The studies here presented represent major advances in the empirical understanding of informal knowledge diffusion mechanisms.

However, a problem common to all the empirical studies presented in this chapter rests in their adherence to the case study. In fact, although quite deep at the level of analysis, these research works are quite narrow in the scope of their investigations; and therefore their findings suffer from lack of generality. This undermines the possibility of drawing general conclusions from such empirical investigations and calls for alternative approaches to the empirical investigation of knowledge diffusion.

We will come back on this issue in the following chapter, where we present a methodological discussion on the core distinction between applied and theoretical models. This investigation will lead us to discuss various agent-based models' validation techniques which could prove to be an alternative and more general approach to empirical investigation on knowledge diffusion patterns.

NOTES

1. As pointed out by Cantner and Graf: '[t]he conscious exchange of technological knowledge between actors can be organized in different types of arrangements. The normative basis for a market organization is a contract between the parties which relies on well defined property rights and actors largely communicate via the price mechanism. Certainly, there are markets for technologies where licences for patents, etc. can be traded . . . The transfer of knowledge can also be organized hierarchically; i.e. within firms where the researcher is obliged to leave the inventions to the employer. Here, the contractual obligations form the basis for a hierarchical structure of coordination' (Cantner and Graf 2006, pp. 463–4).
2. A notable exception is provided by the Community Innovation Survey, an official EU-wide survey that asks business enterprises to report innovation outputs, innovation inputs and, most importantly, sources of knowledge for innovation efforts. Firms are asked to rate the importance of knowledge flows for innovations from a number of sources such as suppliers, other firms in the firm group, customers, universities and so on. Two major disadvantages of these data refer to their nature of being subjective (as they refer to the perceived impact of knowledge flows) and qualitative (as they do not measure the quantity of knowledge flowing in and out the firm).
3. Social network analysis has its historical roots in the disciplines of sociology, social psychology and anthropology, and it focuses on networks' structural description. It represents a distinct research perspective within the social sciences as it is based on the assumption that relationships among interacting units are essential in understanding individual and social dynamics. Therefore it offers theories, models and empirical studies articulated in terms of relational analyses.

In the words of Wasserman and Faust: 'a social network consists of a finite set or sets of actors and the relation or relations defined on them' (1994, p. 20). Such relational ties between actors are channels for the transfer or flow of resources (either material or immaterial) such as knowledge.

4. It is worth noting that the use of the information in patent citations involves some difficulties. One major problem is that patents likely measure a selected form of knowledge increase, since not all innovations are patentable, neither are all patentable innovations chosen to be patented (Crespi et al. 2007).
5. Referring to such a time span allowed the authors to investigate the dynamics of the networks. Specifically, Cantner and Graf (2006) split the sample into two periods of equal length which were subsequently compared. The first period includes all patents disclosed between 1995 and 1997 while the second period covers the years 1999 to 2001.
6. As stated by the authors, the classification into 30 technological classes was obtained employing the International Patent Classification elaborated jointly by the Fraunhofer-Institut für Systemtechnik und Innovationsforschung (FHG-ISI), the Observatoire de Sciences et des Techniques (OST), and the Science and Technology Research Policy Unit of the University of Sussex (SPRU).
7. This database includes information from the German patent office and the European patent office.
8. Note that the database constructed by Meder is strongly unbalanced (due to the large number of zeros present in the dependent variable). As shown by King and Zeng (2001) for strong unbalanced datasets logistic regressions sharply underestimate the probability of rare events and lead to inefficient results. Using the prior correction model allows correction for this shortcoming.
9. Absorptive capacity is here defined as the stock of knowledge accumulated within the firm, embodied in skilled human resources and accrued through in-house learning efforts.
10. This index measures the extent to which technical knowledge originates from a firm to be used by other local firms. The indicator was computed on dichotomous bases (which reflects the presence or absence of such a linkage) and on a valued base (which analyses the value given to each linkage by the knowledge user – ranging from 0 to 3).
11. This index measures the extent to which technical knowledge is acquired by or transferred to a firm from other local firms. This indicator, too, was computed on a dichotomous base and a valued base.
12. This index measures the degree of cognitive interconnectedness of a firm on the basis of its propensity to be in between other firms' knowledge linkages.
13. This index measures the ratio between the knowledge received and that transferred by each firm, and allows distinction among net 'absorber' of knowledge, net 'source' of knowledge and mutual exchanger of knowledge.
14. Firms that have a central position in the network in terms of knowledge transfer to other local firms and that are also strongly connected with external sources of knowledge.
15. Firms that form a central part of the local knowledge system with balanced source-absorber positions within the cluster. They also have relatively strong external links. Although they are less strongly connected to external sources than the TG firms, they behave in a similar way to 'technological gatekeepers' by bridging between external sources and local absorbers of knowledge.
16. Firms that are similar to AMEs in that they are well linked to external knowledge sources and play a relatively balanced source and absorber role within the cluster. However, compared with AMEs, they are less well connected to other firms in the cluster.
17. Firms that have established strong linkages with external sources, but have limited links with the intra-cluster knowledge system. These weak intra-cluster links are primarily inward and absorption-centred.
18. Firms poorly linked at both the local and extra-cluster levels.

19. In their paper the authors refers to knowledge diffusion as ‘knowledge sharing’. For the sake of clarity and homogeneity we will continue to refer to knowledge diffusion patterns.
20. The field work was conducted at the Buijs Group, part of the NatLab, the largest laboratory of Philips Research, and at the Oil and Gas Innovation Research (OGIR), the exploratory research group of Shell Global Solutions.
21. In their paper the authors refer to knowledge sharing as the deployment of knowledge in communication with others. This description is very similar to what we have labelled ‘knowledge diffusion’, that is, an informal and largely uncontrolled flow of knowledge. In what follows we shall stick to the authors’ original notation.

6. Theoretical and applied methodologies of agent-based models

The objective of this chapter is to discuss the approaches of applied and theoretical modelling, to emphasize differences in the nature of modelling enquiry and to suggest that these differences – axioms of modelling methodology – should be recognized as starting points that determine the way the enquiry is planned, carried out and evaluated by modellers.

This is followed by a closely related discussion of validation of agent-based models. Here, validation is considered quite broadly, encompassing both inputs and outputs to the modelling as well as its incorporation into all stages of the model building and analysis. It draws on a diverse review of literature from computer science perspectives, from sociology and economics, and from applied agent-based models.

This discussion will pave the way to the study conducted in the following chapter. The objective is to place the applied model presented in Chapter 7 within the methodological framework introduced, and to give the reader some context for this book outside the specific domain of knowledge and innovation.

THEORETICAL AND APPLIED MODELLING

Broadly speaking there are two streams of investigation that are possible to classify usefully. The first type of enquiry, referred to as foundational agent-based social simulation (Moss 2001), is concerned with the formulation and verification of social theory and the design of agent architectures. Some of the most striking findings about the dynamics of agent-based models were achieved with early foundational models like Schelling's segregation model (Schelling

1978), the Santa Fe Artificial Stock Market (Arthur et al. 1997) or Sugarscape (Epstein and Axtell 1996). The finding that simple rules can generate complex behaviour, which was demonstrated early on, has been a basis for much further interest and subsequent application in a variety of research fields. Such insights of foundational modelling have also led to concerted study of the general properties of complex systems such as non-linearity, emergence, self-organization, and so on.

A second stream of modelling, labelled representational agent-based social simulation, is concerned with the use of multi-agent systems to describe observed social systems (Moss 2001). It has developed as researchers from other domains – as well as decision-makers from industry and the public sector – are finding that agent-based models may be applied to problems of an empirical nature. The problems and questions which they seek to address often have a high level of complexity, and it is in this situation where agent-based modelling is one of the tools which seems to offer some promise.

TYPES OF PROBLEM ADDRESSED WITH AGENT-BASED MODELS

Foundational and representational modelling streams correspond strongly to problems of a theoretical and an applied nature, respectively: in this chapter we use the terms interchangeably. However, as will be presently discussed, the boundary between the two strands of research is not always so clear cut.

Agent-based approaches often integrate social (and economic and ecological) models with physical systems. However, there are important distinctions to draw between physical modelling and social modelling in relation to the normal way that each type of model is confronted with reality – which can be usefully classified into an ‘open’ or a ‘closed’ type of system.

A ‘closed’ system is one typified by a physical science laboratory where all input variables can be controlled or can be precisely measured and are therefore known. Validating a model of a closed system is greatly aided by this ability to perform controlled exploration. One can provide the same inputs to the closed laboratory system as those provided to the (theoretical or simulation) model with a great degree of certainty. The ability to test propositions in the laboratory, by

means of repeated experiments over a large range of parameterizations, can be very helpful in developing scientific knowledge. For the physical world this has resulted in a (well-validated) general theory that has proved to be accurate and has developed cumulatively over the generations.

In the natural and social sciences, on the other hand, it has proved less possible to develop general theory. Knowledge has usually been constrained to local contexts and has not been generally applicable; general theory has not been well validated in relation to real-world observations and experiments. The latter point is a main critique which has been directed towards assumptions of classical economics. In natural science, Oreskes writes (Oreskes et al. 1994), the concept of conventional validation against reality is misleading. Models are bound to be less accurate representations because the systems they purport to represent are 'open' systems – that is, they are real-world systems that do not have a complete, known set of input values. Therefore models will not be able to capture all of the unknowns but must find proxies for these uncertainties, which also means that they will not be able to generate a unique result (or simulation run) comparable with an observation.

The idea of confirmation is introduced (Oreskes et al. 1994) as an imprecise measure for the extent to which a theory or model matches with what is known about the behaviour of an open system. A model or theory should aim for adequacy, it is argued. In empirical social simulation, the target being an open system, CAVES (2007) propose that such confirmation could be provided by activities involving stakeholder validation (see below for further information) of the legitimacy of a model.

There is a related debate which, on the one hand, views agent-based social simulation as a starting point for 'controlled, repeatable experiments to test hypotheses about how the world can be the way it is. Without a laboratory in which to perform these kinds of experiments, there can be no such thing as a bona fide scientific theory of anything' (Casti 1997, p. x). In other words, it is a new approach for social science that calls upon the rigour of natural science methods to generate a qualitatively different type of knowledge. It is an approach which demands, on the other hand, that the 'laboratory' be sufficiently accurate in its depiction of the 'world' that correct inferences can be made from the model. Because of the 'open' nature of social systems and the difficulty of capturing all of the unknowns, this

inference step cannot be guaranteed, and becomes subjective, which will be followed up in the discussion of model validation later in the chapter.

There now follows a discussion of methodological considerations that need to be made in following the theoretical or applied school of investigation. Foundational modelling is undertaken with respect to a target area of social theory or widely held ‘stylized facts’ – empirical regularities derived from a number of studies. Facts may be held in different degrees of certainty, as well as in different degrees of generality. Certainty relates to whether the theory is well established and supported by evidence or whether it is contested, whereas generality relates to whether such statements hold under a wide range of conditions or whether they have in fact been shown to be special cases. All such theories are falsifiable, in principle.

We include logics in the foundational category. For any logical theory, a corresponding simulation model can be built to represent that theory. A number of theorems can then be tested – in the sense of attempting to find a resolution to the theorem. Validation of a theorem is the production of a proof. Conclusions can be drawn about social dynamics based on proving theorems: this is a useful technique and one which does not require additional empirical data. Further foundational knowledge can be derived from analytical proofs, and would also include many existing modelling or simulation studies.

Foundational modelling therefore proceeds towards the specification of a model representing the salient aspects of its target theory. The formalization step involves focusing on aspects of the theory relevant to the study, clarifying these aspects and producing conceptual models, as shown on the left-hand side of Figure 6.1.¹

The conceptual modelling phase terminates with a complete specification of a model which can then be translated into a simulation model, that is, a concrete implementation of the conceptual model ready to be run.² Generally speaking this requires adding assumptions as much as it involves abstracting from a body of theoretical knowledge. This is because a theory is often underspecified in several dimensions as compared to what is desired from an agent-based model, which often requires a lot of design information. So, the left-hand side of Figure 6.1 shows the formalization step as one of adding information in order to specify the theory more concretely in a conceptual model.

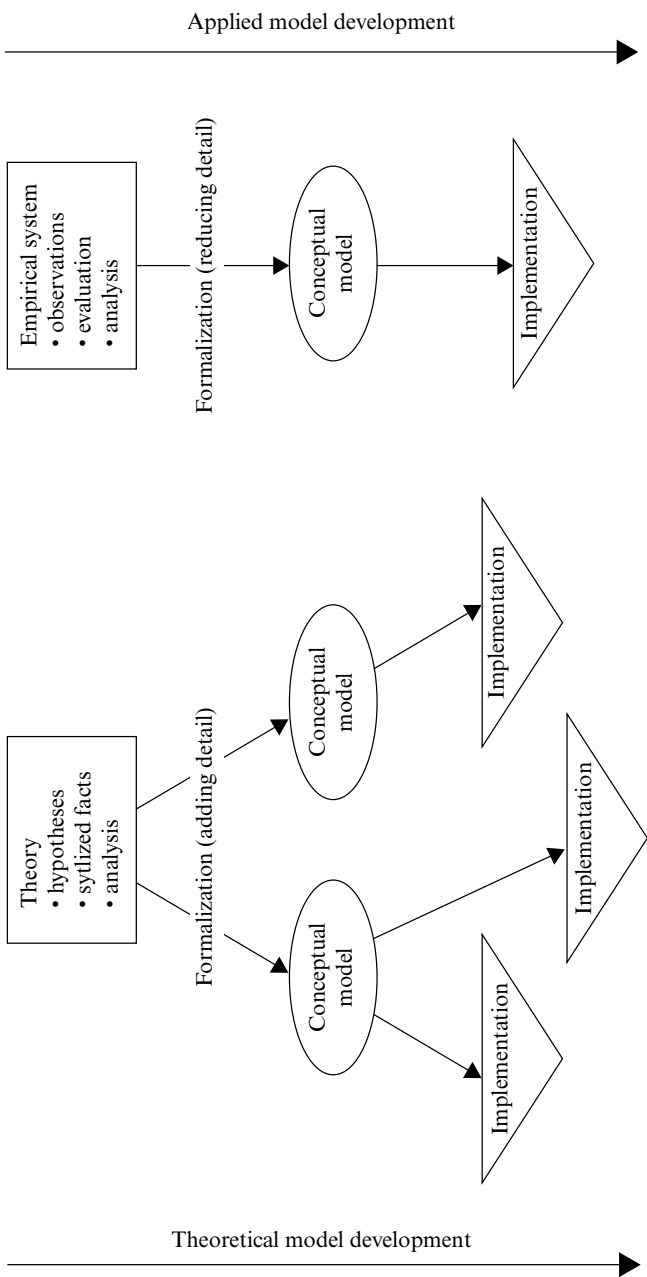


Figure 6.1 Showing target and conceptual model

This step is vital to the actual implementation of the simulation model. Experiments with the theoretical model are carried out with the objective of gaining a better understanding of this type of system and contributing to knowledge by confirming or refuting an existing hypothesis, by inducing that some new (previously unanticipated) social mechanism is at play, or by generating some original facts relating to and expanding the existing general theory.

Now focusing on the empirical modelling where the objective is to apply the methodology to a well-defined study where there is substantive access to field data, it is natural that the abstraction step involves reducing complexity in comparison with the real-world system, and by so doing developing an applied model that, by the nature of its being simpler and more amenable to study, produce some insight into the more complex real-world counterpart. It should be simpler – in terms of being easier to understand the micro and macro behaviour and how these are linked – whilst as accurately as possible representing the relevant aspects of the target.

Generally speaking, the aspects to be represented in the conceptual model are determined by the available information, the intuition of the model designer, and through consultation with any stakeholders who may be influencing the whole study in terms of their problem formulation. As shown in Figure 6.1 the completion of this phase involves reducing the available information to produce a manageable yet sufficiently representative conceptualization (the notion of sufficient simulation, that is, conceptual accuracy, is discussed in the section below on conventional validation).

Empirically focused agent-based simulation aims to address particular social issues that are often quite local and time-specific. It is usually used in conjunction with other methods of field-based research or in conjunction with other decision support tools. However, an agent-based model tends less towards implementation as a management tool. Two of its main strengths when used in an applied context are: (1) bringing forward a common basis of understanding among different parties involved in a study; and (2) exploring the consequences of different possible policy interventions. The latter comes with the proviso that one must remain aware of the large range of uncertainty inherent in complex systems, the lack of accuracy in the results (with respect to the target) and the implications of these facts for model predictability.

What is common across both foundational and applied modelling

is that something about the target is inferred from the behaviour of the formal simulation model. An essential difference is that whereas foundational models relate to largely closed systems, empirical systems are by definition open. It follows from these principles that what can be expected from foundational and empirical models is very different, and that the methodology is different.

The inference step just mentioned is, of course, based on validation of the model. Validation is an important step because it improves modellers' confidence in the results – more precisely, in the strength of the inference from conceptual model results to the target system.

VALIDATION OF AGENT-BASED MODELS

Validation is normally understood as 'determining whether the simulation model is an acceptable representation of the real system given the purpose of the simulation model' (Kleijnen 1999, p. 647). In the agent-based model literature this is becoming an important issue. The reason for the growing interest around validation procedures is strictly related to the nature of the models. In fact, as suggested by many authors, computer-simulated models lack a common methodological protocol with respect both to the way models are presented and to the kind of analysis that are performed (Leombruni et al. 2006). The absence of a shared methodological protocol (which leads to a lack of comparability), coupled with the perceived lack of robustness of agent-based models (Windrum et al. 2007) are, most likely, among the main reasons why mainstream literature (for example orthodox neoclassical economics) is so reluctant to give agent-based simulation models equal dignity to models with closed form solutions. Marks (2007) maintained that better validation would reduce any scepticism about model results and usefulness. The question at stake is how to prove 'that model outcomes reflect persistent (locally generic) aspects of the system under study, rather than a modeller's choice of parameter settings, initial conditions, or software/hardware platform preferences' (CAVES 2007, p. 6).

However, providing a conclusive answer to such questions is a difficult task due to the high degree of complexity of agent-based models (see Chapter 3). This results in the absence of reduced form equations more suitable for empirical validation and in methodological difficulties for developing standardized validation procedures. Whilst it

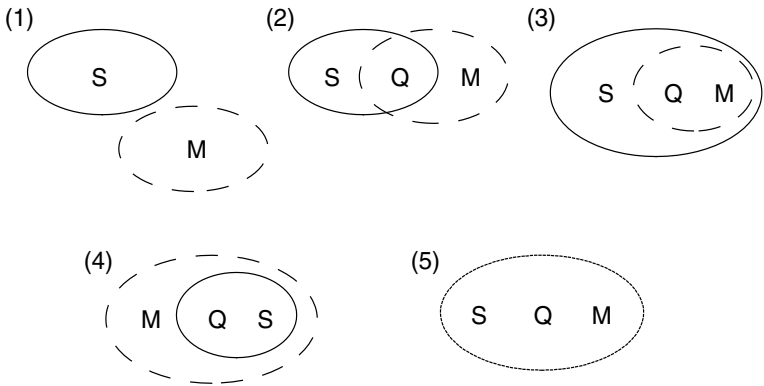
might not be possible to build up a definitive approach to validating agent-based models, researchers have proposed several procedures that modellers can follow to improve the level of confidence placed in the model findings. In what follows we shall attempt to contribute to the debate on validation by suggesting a comprehensive methodological framework of analysis able to embrace various strategies proposed so far in the literature. We will start by surveying the most conventional approaches to validation and subsequently extend the methodological debate to other approaches to validation.

Conventional Approach to Validation

Validation is typically understood as a way of assessing the fit of the model Data Generation Process (mDGP) to the real-world Data Generation Process (rwDGP). In other words, validating an agent-based model requires undertaking several steps which will increase the confidence that the model is a good representation of the real world and will enhance its predictive power. This can be achieved by following two distinct, but not mutually exclusive, approaches: modellers can validate model inputs (for example model parameters) against the real world, and/or validate the output of the model against the historical output of the real world system. The first validation approach is also referred to as micro-level validation (CAVES 2007) or as empirical calibration (Windrum 2007); the second approach applies to the macro level of the analysis and refers to the extent to which modelled results concur with the real world.

Following Marks (2007), we can introduce a general formalization of the conventional approach to macro validation. Let set \mathbf{M} be the exhibited outputs of the model over time and \mathbf{S} be the specific, historical output of the real-world system over time. Let set \mathbf{Q} be the intersection, if any, between the set \mathbf{M} and the set \mathbf{S} , $\mathbf{Q} \equiv \mathbf{M} \cap \mathbf{S}$. We can characterize the model output in five cases:

1. If there is no intersection between \mathbf{M} and \mathbf{S} ($\mathbf{Q} = \emptyset$), then the model is useless.
2. If the intersection \mathbf{Q} is not null, then the model is useful, to some degree. In general, the model will correctly exhibit some real-world system behaviours, will not exhibit other behaviours, and will exhibit some behaviours that have not historically occurred. That is, the model is both incomplete and inaccurate.



Source: Marks (2007).

Figure 6.2 Validity relationships

3. If M is a proper subset of S ($M \subset S$), then all the model's behaviours are correct (match historical behaviours), but the model does not exhibit all behaviour that has historically occurred. The model is accurate but incomplete.
4. If S is a proper subset of M ($S \subset M$), then all historical behaviour is exhibited, but the model will exhibit some behaviours that have not historically occurred. The model is complete but inaccurate.
5. If the set M is equivalent to the set S ($M \Leftrightarrow S$), then the model is complete and accurate.

Hence, according to Marks (2007), the model is incomplete if $S \setminus Q$ is non-null, so that the model does not exhibit all observed historical behaviours. And the model is inaccurate if $M \setminus Q$ is non-null, so that the model exhibits behaviours that are not observed historically. Figure 6.2 illustrates these relationships.

We should, however, note that if on the one hand correlation between real-world data and model output is a necessary condition for accuracy, on the other hand accuracy is not a sufficient condition for causality. 'Simply because the observed data fits the construction of a model does not necessarily mean that the model approximates the reality correctly' (CAVES 2007, p. 10). This

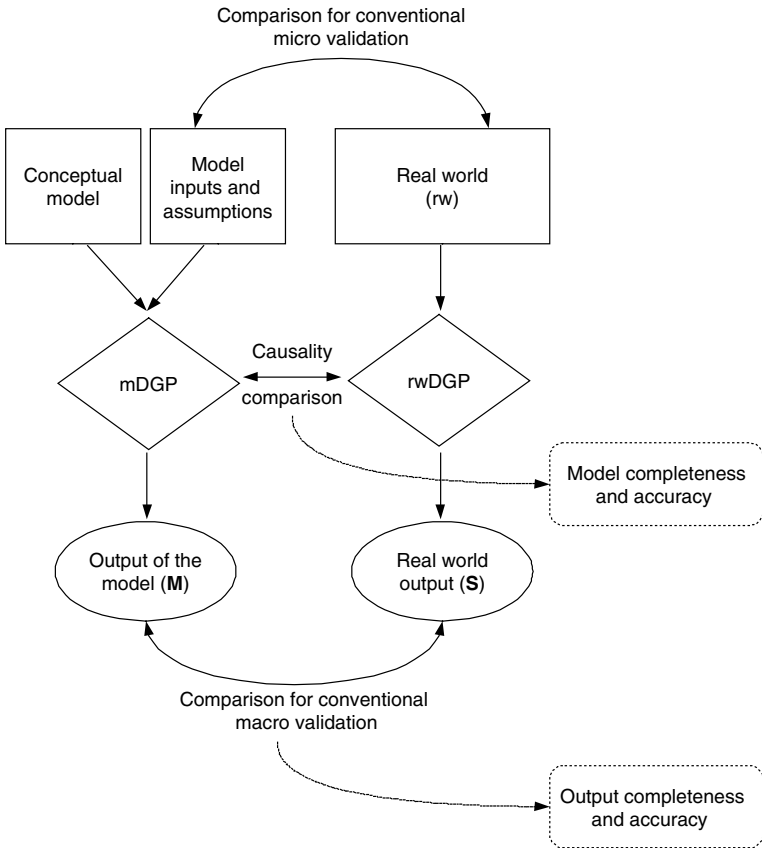
observation raises a general criticism of the formalization just presented as we believe that output accuracy and model accuracy are different concepts which should be kept distinguished. Specifically, the former occurs when historical observed data (real-world data) is a proper subset or is equivalent to the output data generated by the model, whereas the latter occurs when the model Data Generating Process accurately replicates the real world Data Generating Process. We schematically illustrate the different levels of comparison in Figure 6.3.

In Figure 6.3 we distinguish among input comparison (micro validation), output comparison (macro validation) and DGPs comparison. The latter comparison aims at confronting the two data generation processes and, therefore, validating the modelled behavioural rules with the true rules of behaviour. This last level of comparison is indeed the most difficult as the *rwDGP* is extremely hard to investigate; this is especially true when studying complex systems where the *rwDGP* can be completely unobservable. Typically, such comparison goes beyond the conventional approach to validation. We will now present some protocols for output validation which are compatible with the conventional approach just discussed.

Protocols for Validation

Following Windrum et al. (2007), we shall now consider the following three protocols for validation: the indirect calibration approach, the Werker–Brenner approach and the history-friendly approach.³

The first approach, the indirect calibration, is quite straightforward: it first performs validation, and then indirectly calibrates the model by making use of the parameters that are consistent with output validation. The empirical validation process is constituted by four pragmatic steps: first, the modeller should identify a set of stylized facts that he or she is interested in reproducing and/or explaining. In the second step, the modeller should create the model in a way that the *mDGP* is kept as close as possible to empirical and experimental evidence about behaviour and interactions. This process implies drawing together all possible data about the underlying principles that inform real-world behaviours so that the model's context is not a too unrealistic one.⁴ This empirical evidence on the stylized facts should then be used to limit the parameters' space as well as to redefine the initial conditions in case the model turns out



Source: Based on CAVES (2007).

Figure 6.3 Comparison for validation

to be non-ergodic.⁵ Following this third stage, the modeller, in the final step, should aim at understanding in a more thorough way the causal relations characterizing the stylized facts in question and/or investigate whether new stylized facts (different from those observed in the first step of the validation process) came to the fore. These new stylized facts, under certain circumstances, can be validated by the modeller *ex post* by further investigating the subspace of parameters that have proven to be resistant to the third step, that is, those consistent with the stylized facts of interest.

The second approach to empirical calibration of agent-based models is the Werker–Brenner procedure, which is made up of three steps (Werker and Brenner 2004). While the first two steps are similar to the calibration exercises just discussed, the third step presents a novelty. The main difference lies in the fact that here the empirical parameters are chosen directly to calibrate the model. More specifically, step one refers to existing empirical evidence in order to calibrate the model's initial conditions and the ranges of its parameters, which should be specified for parameters for which there is little or no reliable data. In step two the modeller performs the empirical validation of the outputs for each of the model specifications defined in step one. This should help to reduce the plausible range of values within the initial dimension space. According to Werker and Brenner, one way of actually carrying out this output validation could be using the Bayesian inference procedures, where each model specification is assigned a likelihood of being accepted, having in mind the percentage of 'theoretical realizations' that are compatible with each 'empirical realization'. In this way, it is possible to limit further the initial set of model specifications (parameter values) as the modeller keeps only those parameter values which received the highest likelihood of being accepted by the current empirical realization. On the other hand, model specifications that are not compatible with currently known data are set aside.

Following this comes step three which entails a calibration process which draws on those sets of models which survived the second step. This step, which might involve consulting expert testimony from historians, is called by Werker and Brenner 'methodological abduction'. Now, the underlying structural model is identified from the shared properties and characteristics of the 'surviving' models. The authors argue that: 'these [shared] characteristics can be expected to hold also for the real systems (given the development of the model has not included any crucial and false premises)' (Werker and Brenner 2004, p. 13).

The third approach to validation is the history-friendly one which, like the calibration approaches discussed above, aims at creating a model more in line with the empirical evidence. However, it is different from other approaches in that it makes reference to the specific historical case studies of an industry to model parameters, agents' interactions and agents' decision rules. In other words, it is an approach which uses specific historical outlines in order to calibrate the agent-based model.

According to this approach, a ‘good’ model is one that can produce various stylized facts observed in an industry (Malerba et al. 1999; Malerba and Orsenigo 2001). Malerba et al. (1999) outlined the approach and then applied it, examining the transition in the computer industry from mainframes to desktop PCs. Malerba and Orsenigo (2001) subsequently applied the approach to the role played by biotech firms in the pharmaceutical industry. This approach entails the construction of industry-based, agent-based models, where detailed empirical data on an industry are required for the model building, analysis and validation. Hence, models are based – in both the building and the validation phases – on a wide range of available data, which are gathered through empirical studies as well as anecdotal evidence or histories written about the industry in question. These data assist the modeller in specifying the model’s agents (their behaviour, decision rules and interactions), and the environment in which they operate. They also facilitate the identification of initial conditions and parameter values of key variables likely to produce the observed history.

Finally, in the model validation phase, the gathered data can be used for comparison of the model output (the ‘simulated trace history’) with the ‘actual’ history of the industry. It is precisely in this dimension that the history-friendly approach distinguishes itself from other approaches: by giving weight to the use of historical case studies to guide the specification of agents and environment, and to identify possible key parameters. The followers of the history-friendly approach claim that it is possible to reach a correct set of structural assumptions, parameter settings and initial conditions through a process of backward induction. As put by the authors, such a correct set of ‘history-replicating parameters’ would facilitate the conduction of a sensitivity analysis to establish whether ‘history divergent’ results are possible.

Alternative Approaches to Validation

The conventional validation approach discussed above, aims for what could be labelled ‘replicative validity’: the model matches data already acquired from the real system (Zeigler 1985 [1976]). However, we can augment the conventional approach to validation by considering two other types of validation: ‘predictive validity’: the model matches data before data are acquired from the real

system; and ‘structural validity’: the model ‘not only reproduces the observed real system behaviour, but truly reflects the way in which the real system operates to produce this behaviour’ (Zeigler 1985 [1976], p. 5).

Basically, predictive validity requires *ex post* replicative validity, whereas structural validity requires validation of the mDGP, which can only be obtained by comparing it with the rwDGP. In practical terms structural validity necessitates, first and foremost, an effective model design. In fact, as maintained by various authors (CAVES 2007), this is a crucial step which might include in built validation techniques.

Recently, Edmonds and Moss (2005) raised the question of whether a model design should be kept as simple as possible (the so-called KISS approach – Keep It Simple, Stupid) or as descriptive as possible (what the authors have labelled KIDS – Keep It Descriptive, Stupid). This question intrinsically relates to the issue of validation. In fact, it has been argued that: ‘KIDS models, because of their descriptive content, are more amenable to scrutiny and criticism by stakeholders and experts than KISS models, which too easily become enigmatic and obscure. This attribute of KIDS models also makes them easier to validate and less capable of “garbage in garbage out” errors’ (CAVES 2007, p. 8).

A different approach to model design is the so-called TAPAS (Take A Previous Model and Add Something) method which is potentially compatible with both KISS and KIDS. Such an approach enables the docking technique of validation which involves the conceptual alignment of models and can be used to check whether an agent-based model (ABM) is a special case of another model (Axtell et al. 1996). ‘Docking probably has the most potential to provide ABM with a standardized validation technique’ (CAVES 2007, p. 9); however, it does not replace the need for validation against the real-world environment as the conventional understanding of validation checks whether the model is a good representative of the reality (CAVES 2007, p. 20).

VALIDATION BY COMPARING MODELS

The previous sections have detailed a number of approaches to model validation, demonstrating the various steps entailed: for

example, correlation between model outputs and historical outcomes, observation of real-world data-generating processes, conceptual accuracy, and induction of stylized facts from empirical and experimental evidence, not all of which involve direct comparison with empirical data. A number of additional validation procedures involve comparison among models.

Figure 6.4 expands on the model development process described earlier by illustrating a further step where more than one model is considered together, and can hence improve the validation. Recall that the formalization step involves specifying the theory or empirical knowledge more precisely into a set of rules and processes suitable to be translated into a computer program, that is, a simulation model.

Figure 6.4 shows two implementations of the same conceptual model. It might be desirable to use different architectures or programming platforms to study the model, for example differences in the speed of simulation or ease of programming might be an important factor, whilst it is also clearly valuable to test several models, thus reducing the possibilities of error. This type of comparison, replication, involves reviewing both the conceptual model and the alternative implementation of that model, and checking that inputs and outputs are the same (Axtell et al. 1996; Edmonds and Hales 2003). This step therefore incorporates verification, one of its objectives being to check that the model is behaving according to the design, as well as validation. Validation is implicit in replication because the comparison may lead to respecification of the conceptual model in order to bring it into closer convergence with the target system. In other words, differences among implementations might reveal assumptions that have not been made explicit, and any new assumptions revealed in this way may improve confidence in the model.

The figure also shows that several conceptual or aligning models may be developed from the same area of theory. Docking is the process of comparing two simulation models when they are close enough to show the same features and behaviours: in other words, when the two respective conceptual models overlap significantly. The experiment will investigate a subset of the space of input assumptions that overlap, and will compare outputs of the two models for consistency (for example Axtell et al. 1996). The objective, as discussed in the above section, is to validate both models.

The thesis of this chapter is that approaches to validation

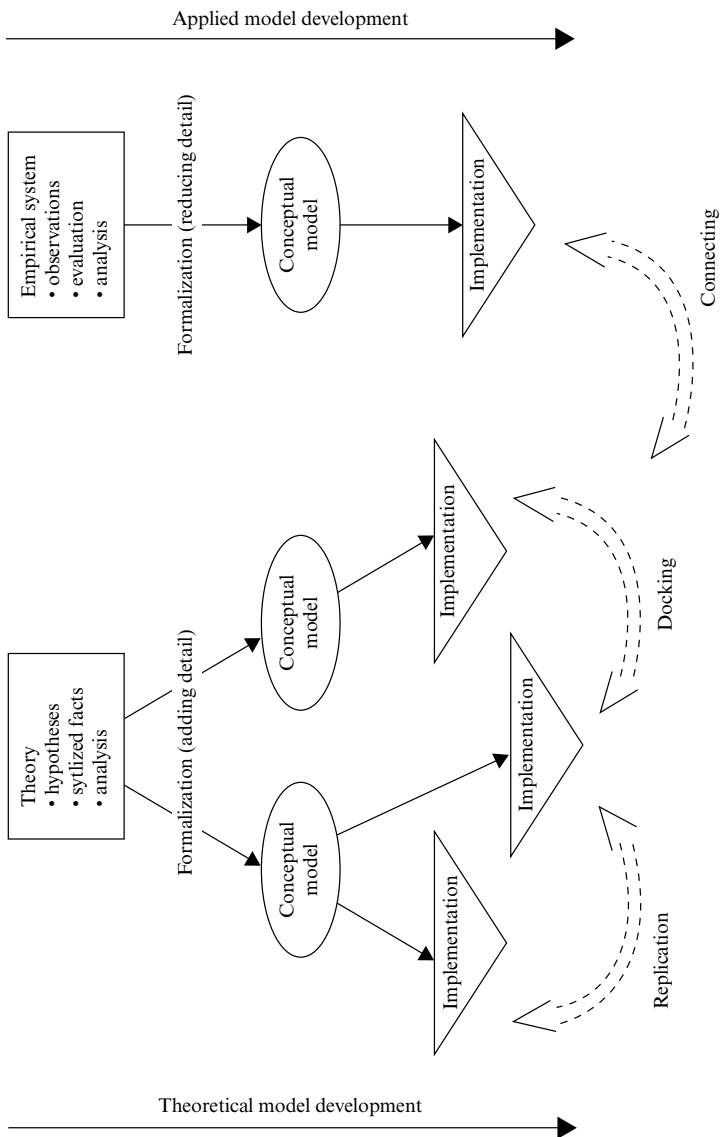


Figure 6.4 Showing validation by comparison of models

necessarily differ according to the type of modelling undertaken, whether empirical or foundational, and that this necessity is due to the nature of the target system. It is argued that empirical modelling addresses the behaviour of open systems, where models are impossible to validate in the conventional sense of obtaining an accurate representation. In this case, validation must be aimed at addressing critiques that relate only to the purpose of the model and the particular focus aspects of the target, with an understanding that such an enterprise will never be complete, and that all input factors are not known or are uncertain in behaviour.

Foundational modelling, on the other hand, has a different set of considerations. First, it is more feasible for conventional methods of validation to be applied, as the areas of uncertainty are here less of a constraining factor. Second, the whole parameter space can likely be investigated in a way that is not possible when addressing real-world complexity.

Validation involves comparing simulation results with earlier analyses and knowledge of the area. It is also more likely to reach a definitive conclusion in relation to previous findings, as opposed to applied modelling which has an iterative character (and has the subjective concept of confirmation as discussed earlier). Output validation corresponds with the more conventional scientific notion of the term (macro validation), as discussed earlier (see Figure 6.3).

Therefore, it is natural to expect differences in the methods used. These factors contribute to differences in methodological protocol (and the lack of common standards, as discussed earlier). Different methods that have been used can also be divided roughly along these lines, that is: (1) theoretical model validation involves conventional methods supplemented by computer simulation – laboratory experiments, role-playing games, theorem proving, sensitivity analysis; (2) applied model validation implies social research, statistical and ‘real’ (non-orthodox, for example evolutionary economic as opposed to neoclassical economic) techniques including model calibration and stakeholder validation.

However, having made clear these differences, it is worth pointing out that the two approaches are often closely linked in practice. Figure 6.4 shows a third type of comparison among agent-based model implementations in which there is an interesting crossover point in validation. The figure shows the divergence in model development pathways which follow from these two streams, and it shows

where they might explicitly cross over – what we have labelled ‘connecting’ models.

The research project may commence with development of an abstract model of an archetypal social situation, or it may commence with development of a detailed model of a specific empirical system. In the first case, the research project may commence with a hypothesis about how a particular social process is thought to occur. This theory may be enriched with a limited number of generalized facts drawn from the literature or from the results of other theoretical studies. The model will be developed, along with a choice of parameters, to capture a wide domain and the conclusions will be about general behaviours of this kind of system. The study may then be followed by a more empirical type of investigation, where the model is applied to a given situation which is local to a region, or to a period in time for example, or in any case to a specific instance of the phenomenon being investigated. The advantage of this approach is that the general model can be used to support and guide the empirical data-gathering stages. The danger is that the model may be misguided in its initial assumptions and does not, therefore, have utility in the empirical study.

The second case starts with a case study-centric investigation, collecting data and involving location-specific expertise. This is followed by development of a more generic model aiming to abstract from the former model and explore the more systematic outcomes of classes of models.

In either case, this would involve not only selecting a different grain of analysis; it also means selecting a number of different assumptions. If selecting from an alternative theoretical or empirical focus, it also encompasses methods that link foundational and empirical modelling.

It is possible to apply one or more validation methods in order to switch from a theoretical to an applied context using, for example, input validation, and this is what will be demonstrated in Chapter 7. Throughout the presentation, we shall be concerned with the connection between these two models. Connecting models is the activity of comparing the structure and behaviour of an applied model with that of a theoretical model. The aim is to validate models in relation to other models, and to increase their scope and utility in both theoretical and applied contexts.

In this section it was argued that the way in which validation

protocol, approach and activities are chosen differs largely across the theoretical–applied axis. Model calibration was discussed in depth above. Now we shall briefly introduce the other methods mentioned in the text.

FURTHER VALIDATION METHODS

Sensitivity analysis is a method for understanding the relationship between the model inputs, in terms of parameter choices or assumptions data, and the model outputs. After determining which input variables are to be analysed, simulations are carried out and the sensitivity of the outputs to variations in the inputs are measured.

Laboratory experiments are used, most notably in economics, to make a simplification of complex real-world processes and put them under rigorous scientific scrutiny. The experimental approach involves constructing an artificial system in the laboratory where all of the inputs can be carefully controlled and the interactions monitored. This can be combined with a simulation experiment approach; that is, the laboratory experiment is compared with a simulation experiment, and experimental findings then lead to inferences about the nature of the artificial system. In this scenario, the target of the simulation model is the artificial (observed) system. An important distinction needs to be made between the artificial system, which is a closed system, and the kind of real-world system addressed by an empirical model, which is not closed.

Role-playing games are a related technique often used in participatory research to elicit local knowledge or to co-learn on certain issues with stakeholders. An exercise in which initial game rules are similarly prescribed can help to inform on which individual strategies players are following and how they are interacting, especially in terms of group dynamics. Follow-up discussions can be of a reflective nature and may in fact be open-ended. This leads to a constructivist approach: the model underlying the game can be adapted as a shared representation among different stakeholders.

Stakeholder validation focuses on feedback processes with the stakeholders, that is, it provides the opportunity to review either the decision rules or the model itself through contact with participants or key informants. This could involve, for example, checking the validity of model assumptions according to stakeholders' expertise, or

checking the outputs of simulations against what has been observed by stakeholders or according to what they believe to be plausible. It is a way of comparing mDGP with rwDGP with those not familiar and knowledgeable about the latter.

The idea of theorem proving is to relate the content of a simulation with a theory that the simulation model corresponds to. If it is possible to generate all simulation trajectories and show that a particular theorem holds true in all simulation model results, this can be considered a proof of the theorem. In other words, a semantic proof of certain regularity in a simulation can be implemented by generating all trajectories and observing whether the regularity holds in all of them (Terán 2001). This can only be attempted for models with a highly restricted number of inputs or decision rules or agents, due to the high computational requirements for carrying out the simulations, and so it is of limited use for generating knowledge about empirical ABMs (which are computationally intensive in particular).

CONCLUSIONS

In this chapter the approaches of theoretical and applied modelling have been discussed. It was argued that the development and testing of theory falls into the foundational modelling classification whereas the application of models to real-world cases is described by representational agent-based modelling. Figure 6.1 showed the main stages in the model-building process, which begins with the definition of a target system. The rest of the chapter then concentrated on the role of model validation, which aims to ensure good representation of the target. We discussed both the conventional meaning of validation in physical science as well as how this can be approached when dealing with uncertainties in natural and social science, and particularly in the relatively new area of agent-based social simulation.

Several protocols for validation were discussed in some detail, from which it could be concluded that validation of model data-generating processes (mDGP) is perhaps the most challenging area. The next section introduced three ways of comparing model implementations, methods which are clearly open to the comparison of mDGP. Replication, docking and connecting models can also contribute to validation, it was argued, because they can result in

respecification of the conceptual model. Here, we made several additional remarks on the method which we labelled ‘connecting’ models where, interestingly, applied models are compared with theoretical models in order to generate further insights.

Validation is important to establish confidence in the use of models to address and analyse real-life problems. However we also argued that it should be considered when addressing foundational questions of a socio-economic nature, in other words when the target of the modelling is an area of social or economic theory. This could be done in any of the ways we have been discussing, that is, by comparison with other models (mDGP) to improve the conceptual model, input validation to increase the relevance of the model to a particular case, output validation (macro validation) to test for the reproduction of stylized facts. Then, we suggested specific methods that we would consider using for theoretical and applied model validation and described these in more detail.

Having discussed the relevant methods of validation and how they are selected according to the modelling objectives, the following chapter will illustrate how this might be done by returning to the model presented in Chapter 4 and calibrating this model with input data from a case study.

NOTES

1. Not all conceptual models, of course, are formal models; in the present context formal conceptual models are the ones of interest.
2. This end-game involves ironing out inconsistencies and unclear aspects of the conceptual model so that the logic is complete.
3. For an extensive discussion of these three approaches and their limitations, see Windrum et al. (2007) and Windrum (2007).
4. As mentioned earlier, this second step can be difficult to establish when dealing with complex systems because they are open systems.
5. In fact, an ergodic system that evolves for a long time ‘forgets’ its initial state.

7. Validating the model of knowledge diffusion and innovation

Elaborating on the methodological discussion developed in Chapter 6, we shall now attempt to validate the knowledge diffusion model developed in Chapter 4. Specifically, we will make use of a micro-validation approach, ensuring that the structural conditions and geographical arrangements incorporated into the model capture the salient aspects of the empirical system under investigation. We will do so by means of an indirect calibration to validate model inputs (that is, initial parameterization) against the real world (that is, data collected from case study).¹ This calibration exercise should allow an increase in the adherence of the model to a real case study and should provide an example of how validation of an applied model can be conducted.

The chapter is structured as follows: first, two applied models on knowledge diffusion are reviewed and the employed validation technique is discussed; then, the case study used to calibrate the agent-based model developed in Chapter 4 is introduced; hence, validation is implemented, results are examined and a set of policy actions to enhance the innovative performance in the system are considered. Finally, some concluding remarks both on the relevance of the validating methodology as well as on the achievements obtained with this validation exercise are presented.

A BRIEF REVIEW OF VALIDATING KNOWLEDGE DIFFUSION MODELS

In this section we briefly present two applied models of knowledge diffusion which have been validated using two different methodologies. The first model (Morone and Taylor 2004b) here

discussed uses calibration as a validating tool and attempts to investigate how knowledge diffuses among individuals endowed with different levels of education by means of face-to-face interactions. Primarily, the authors focus upon the processes of knowledge diffusion and learning, and the emerging network characteristics that result from these interactions. The calibrated simulation model makes use of socio-geographical data from the Greater Santiago de Chile case study. This model could be considered a calibrated extension of Morone and Taylor (2004a), discussed in Chapter 3, as it departs from similar assumptions but develops a more complex learning structure.

Specifically, the authors consider a population of 232 heterogeneous agents distributed over a grid that resembles the geographical configuration of the metropolitan area of Greater Santiago de Chile. The grid is divided into 34 *comuna* (districts), each corresponding to a defined district of Santiago and thus having different dimensions and different population densities.² Each agent is initially assigned a district and then allocated to a cell at random within that district. Neighbourhoods, which may contain cells from several districts, are constructed according to the von Neumann definition and the initial local network is created by connecting an agent with all other agents located within her/his neighbourhood. It also defines a cyber network as the ideal network connecting all those agents who have access to the Internet. The cyber network generates a second system which has no geographical dimension but connects agents located in far distant regions through information and communication technology (ICT) support. As opposed to the local network, all agents who have access to the Internet are neighbours within the cyber network (that is, the visibility is equal to the size of the system).

Each agent has a list of acquaintances which includes the members of the local network and the cyber network. However, the resulting network structure changes over time as each agent can learn of the existence of other agents through interactions with her/his local neighbours (that is, she/he can be introduced to the acquaintances of her/his acquaintances). Moreover, agents can stop interacting with some of their acquaintances if the connection does not tend to result in useful interactions. Hence, the number of acquaintances changes over time, but does not necessarily increase over time.

Each agent is also endowed with a complex cognitive structure –

that is, the cognitive map – which contains information on the level and kind of knowledge he/she possesses. The structure of the cognitive map is that of a tree which resembles, although in a simpler way, the firms' skills universe depicted in Chapter 4. The data which are used to calibrate and initialize the model are a subsample of the 1998 edition of the *Encuesta de Ocupación y Desocupación* and provide information on the location of agents in the grid, their initial level of knowledge and specialization (that is, measured as years of schooling and kind of school attended), and whether they have access or not to ICT facilities (that is, whether they belong or not to the cyber network).

The basic objective of the simulation experiments is to test whether knowledge diffuses homogeneously throughout the population or whether it follows some biased path generating divergence and inequality. By tuning the initial parameters the authors were able to simulate the impact of a policy intervention and, therefore, forecast the impact of policy action aiming at increasing the minimum level of schooling, or enhancing the diffusion of information and communication technologies.

The results of the simulation exercise suggested that under the initial conditions observed in Greater Santiago de Chile (that is, a high level of knowledge inequality) there is a high risk of exclusion for those people initially endowed with a low level of education. Moreover, studying the spatial dimension of the exclusion process, the authors found that knowledge inequality is more likely to increase if an initial situation of a low level of knowledge is coupled with geographical exclusion. In other words, those people who start with a relatively high level of knowledge (that is, schooling) will always be able to gain from informal face-to-face learning, while those with a lower level of schooling might be trapped in a lower knowledge equilibrium if their low level of knowledge is cumulated with geographical exclusion.

As mentioned above, these findings allow the authors to forecast the impact of policy action aimed at reducing the knowledge gap. Specifically, the authors identified two possible policy actions which could help to curb knowledge inequality: the policy-maker could increase the level of education of more backward and marginalized people, and/or reduce the geographical gap between centre and periphery. This latter policy could be implemented through the development of infrastructure bridging the centre-periphery distance,

as well as through investments in ICT especially concentrated in peripheral areas.

In a second paper Morone et al. (2007) employed a different technique to validate a similar agent-based model of knowledge diffusion. Again, the authors considered a population of heterogeneous agents endowed with a cognitive map representing their knowledge structure, in which exchange of knowledge is by means of face-to-face interactions. As opposed to the previous model, in this paper agents were allocated in a one-dimensional wrapped grid (that is, a ring) and three static network structures were considered: a regular network, a random network and a small world network.³

The authors aimed at defining the main factors which influence the speed and the distribution of knowledge diffusion within a closed network. They identified four fundamental factors: (1) the learning strategies adopted by heterogeneous agents; (2) the network architecture within which the interaction takes place; (3) the geographical distribution of agents and their relative initial levels of knowledge; (4) the network size. Validating the model by means of a laboratory experiment, the authors attempted to single out the impact of each of these factors on learning dynamics.

The experiment was run at the University of Bari in the ESSA laboratory, and 14 players took part in the experiment. Each experimental player was initially endowed with a cognitive map representing his/her cognitive structure and initial knowledge level. The aim of the game for each player was to increase his/her own level of knowledge by means of face-to-face knowledge exchanges. Each experiment lasted 100 time steps and during each time step each player had the opportunity to acquire a bit of knowledge by interacting with one of his/her neighbours. The setting for knowledge flows was that of a gift economy in which players exchanged knowledge freely.

From a methodological point of view, combining a laboratory experiment with an agent-based model allowed the authors to perform both a macro-level validation (that is, comparing the output of the model against the 'real-world' experimental results) and a more articulated validation of the model Data Generation Process (mDGP) against the real-world Data Generation Process (rwDGP). This second validating step was conducted by investigating the learning strategies adopted in the lab and developing a simulated strategy able to replicate those strategies. Moving along

the footsteps of evolutionary economists like Nelson and Winter (1982), Silverberg et al. (1988) and Dosi et al. (1995), the authors argued that they developed a sort of history-friendly model which was labelled an experimentally friendly model (refer to Chapter 6 for further discussion on experimental studies as a means of agent-based model validation).

Investigating the experimental results, it emerged that players followed, in the vast majority of cases (almost 60 per cent), a ‘width-first learning strategy’ – that is, a learning strategy in which agents preferred broadening the scope of their knowledge rather than acquiring more specialized knowledge. Moreover, letting simulation agents replicate such a strategy led to a model able to mimic quite closely the experimental learning pattern. This finding was then compared with those obtained using a set of alternative simulated learning strategies and showed that the output validation was best achieved when this component of the mDGP resembled what we had observed in the laboratory, that is, the width-first learning strategy.

Further results were obtained by examining those factors which affect knowledge flows. First and foremost it was shown that learning dynamics is affected by the learning opportunities provided to each agent in the network. By ‘learning opportunities’ the authors mean the chances each agent has to interact with knowledgeable agents. From this finding it followed that a particular geographical distribution of agents (endowed with different knowledge) could substantially affect learning dynamics. In order to test independently the effect upon learning dynamics of the network structure and of the geographical distribution of agents, the authors ran batches of simulations for each network’s structure, while allocating the agents in different ways for each simulation. Then, they computed the average performance of each network – hence clearing out the geographical effect. Once they had been corrected for any possible geographical bias, the authors could conclude that small world networks do perform better than regular networks, but consistently underperform when compared with random networks.

The studies here discussed provide two alternative approaches to agent-based model validation and, according to the methodological discussion developed in Chapter 6, serve different purposes: the calibrated model study being designed to forecast the impact of policy action (which is simulated by changing the initial conditions tuning key parameter values), and the experimental study being

more suitable for output validation and DGPs comparisons. In what follows we shall present a case study on organic production in a backward area of southern Italy which will be the base for the calibration exercise of the model of knowledge diffusion and innovation presented in Chapter 4.

ON THE CASE STUDY

The model developed in Chapter 4 will now be calibrated using primary data collected in a backward area located in the south of Italy (the province of Foggia) on a group of firms operating in the food sector and, specifically, involved in the production of organic food. The food sector has always been regarded as a low-technology industry, having been associated with low-tech transformation of agricultural products. In comparison with other major industrial sectors, research and development activities in the food industry are of minor weight. All the same, the food industry has gone through dramatic changes during the past few decades.

First, the food industry, like most of the manufacturing sector, experiences increased competition both on the domestic and on the international markets. Second, it has been losing the confidence of the general public due to a series of severe food safety shocks, such as the 'mad cow' (BSE) or foot and mouth diseases in the UK or the dioxin in chicken crisis in Belgium. Finally, environmental and cultural concerns have entered the food debate, directing consumers' attention towards issues of long-term ecological sustainability as well as animal rights (Boudouropoulos and Arvanitoyannis 2000). In this sense, the fundamental challenge that the agricultural sector and the food supply chain have been facing over the last few years (and indeed are still facing) is to move from a supply-oriented strategy to a demand-led one, driven not only by economic considerations, but also by social, cultural, ecological and other values which reflect changing consumer preferences and new legal and regulatory frameworks.

In response to these changes, the food industry found itself in a situation where it was forced to introduce innovations as a response to rising pressures from two fronts. On the one hand, food enterprises needed to keep up with stricter regulation covering food safety, food quality and environmental standards. On the other hand, pressure

also arose from various stakeholders (for example environmental non-governmental organizations or the general public) to go beyond these statutory regulations. As a result, food standards and labels (such as organic food) were developed in order to identify companies that implemented strategies that went in this direction.

Avermaete and Viaene identified three innovation strategies often adopted as a response to such pressures: (1) food safety and quality systems; (2) environmental management strategies; and (3) labelling. 'In contrast with conventional innovation, these strategies are based on innovations for which the procedures are set by an external party. At the same time, the strategies go far beyond technological innovations. Information, communication and networking play a key role for successful implementation of the three mentioned strategies' (Avermaete and Viaene, 2002, p. 3).

The processing of organic products represents a type of innovation which encompasses the three strategies mentioned above. It differs from traditional innovation strategies since, rather than an incremental product innovation or basic process innovations, organic producers actually reposition themselves and their whole product line on the principles of organic production and at the same time keep their traditional product line (Grunert et al. 1997). Hence, switching to organic production is in itself an innovation which results in the production of an output that better responds to changes in consumers' tastes. However, once switched to organic production, firms will constantly have to comply with externally defined standards and product characteristics; this will require constant innovating and learning capabilities. Such capabilities rely largely on firms' aptitude to optimize communication within and among themselves in order to guarantee the achievement of specific food standards.

Along this line of reasoning, our applied study aims to investigate how relevant knowledge diffusion is for organic production. Using the Foggia organic food database will allow us to investigate and model the potential innovative capability of this cluster of firms. In the area of Foggia there are 120 organic industrial firms out of which we chose a cluster of 32 units selected with the focus group technique.⁴ This technique provides qualitative information on a specific theme by playing on the interaction and confrontation of points of view expressed by participants in a discussion conducted by a facilitator (European Commission, 1999). In this case study,

participants of the focus group belonged to local public institutions, research centres, entrepreneurial associations and certification agencies (that is, quality control agencies), from whose interaction we obtained a draft of the organic firms' network that was checked and corrected during the direct survey. Note that, although relevant,⁵ in this study we are not considering extra-cluster knowledge relations nor firms–institutions knowledge flows, as we focus our attention solely on intra-cluster knowledge flows which occur among firms. In Figure 7.1 we report the geographical distribution of the 32 firms. A list of the corresponding names, the province to which they belong and knowledge base⁶ is reported in Table 7.1.

The questionnaire, which was submitted with face-to-face interviews, was structured in two parts. The first part aimed at gathering general information on the characteristics, location and knowledge base of the firm. The second part aimed at collecting information on relations and, more precisely, on the existence or not of ties, their nature and, in the case of communicative relations, the kind of information exchanged.

These data were used in an earlier work (Morone et al. 2006 – see Chapter 5) where we showed how organic food producers proved not to be very efficient in transferring knowledge among involved agents. However, the underlying network's architecture showed potential for knowledge transfers. Implementing a calibrated model of innovation fostered by interactive learning will allow us to investigate the true potential of this network of firms and, eventually, prescribe some policy measures to enhance the knowledge diffusion and ability to innovate of such a network.

RESULTS OF THE SIMULATION MODEL

As mentioned earlier, we present here the results obtained using a calibrated version of the model discussed in Chapter 4. For this model, we also used the JAVA platform with the RePast (Recursive Porus Agent Simulation Toolkit, North et al. 2006) libraries for implementing the model and JUNG libraries (Java Universal Network/Graph Framework, 2007; O'Madadhain et al., 2005) for analysis of the networks data. The model architecture is identical to that described in the earlier chapter. However, some of the data collected from the case study described above will be used to parameterize the model.




Figure 7.1 Firms' geographical distribution

Table 7.1 Firms' knowledge base

No.	Interviewed firms	Location	Knowledge base of the firm		
			Low	Medium	High
1	ANGARANO SRL	Cerignola		x	
2	BIOFACTORY SAS	Cerignola	x		
3	OLEIFICIO SAN SALVATORE SNC	Stornara		x	
4	DONNALISA SRL	San Severo	x		
5	OLIVETO BELMONTE DI BALDASSARRE ANGELICA	San Severo			x
6	AZIENDA AGRICOLA DOTT. CACCAVO	Foggia		x	
7	CIPAM SCARL	Foggia		x	
8	COSEME SRL	Foggia		x	
9	SANTO STEFANO SRL	Cerignola	x		
10	B AND B SRL	Stornara		x	
11	NAPPI SRL	Stornara		x	
12	ORTODAUNIA SRL	Stornara	x		
13	CI.SEME SNC	Stornarella	x		
14	MOLITORIA NUSCO SRL	Ascoli Satriano		x	
15	SARACINO MIRIAM	Ascoli Satriano		x	
16	OLEIFICIO LE FASCINE SRL	San Severo		x	
17	L'AGRICOLA PAGLIONE	Lucera			x
18	JOLLY SGAMBARO SRL	Cerignola		x	
19	IL PARCO DI CASTIGLIEGO MARIA	San Giovanni R.		x	
20	FRANCO LA DOGANA & C. SAS	Orta nova			x
21	SOTTO LE STELLE DI PALLADINO RACHELE	San Giovanni R.			x
22	D'ARIES ANTONIO	Lucera		x	
23	LA QUERCIA SCARL	Foggia		x	
24	LA NUOVA ARPI SCARL	Foggia		x	
25	SPIAVENTO SRL	San Severo		x	
26	CANTINE VINARIS SAS	San Severo		x	
27	SANTACROCE	Deliceto	x		

Table 7.1 (continued)

No.	Interviewed firms	Location	Knowledge base of the firm		
			Low	Medium	High
28	PARADISO TOMMASO & C. SAS	Cerignola	x		
29	EMMAUS SOC. COOP.	Foggia		x	
30	EUROAGRO ALI-MENTARE	Ascoli Satriano		x	
31	DI TUCCIO RAFFAELE	Cerignola			x
32	PUGLISSIMA SRL	Cerignola			x

Specifically, we will use field data to define the network within which firms operate and we will use information collected about firms' knowledge bases in order to define the initial skills profile of each firm. In Table 7.1 we classified firms into three groups according to the amount of training activity undertaken. We have firms with a high knowledge base (HKB), firms with a medium knowledge base (MKB) and firms with a low knowledge base (LKB).

Except for the calibrated model settings defining the number of agents, the knowledge bases and the network structure,⁷ we used the same fixed parameters as those employed in Chapter 4. In other words we have the following: number of time steps = 100, number of agents = 32, number of radical innovations = 60, number of incremental innovations = 60, number of skills per innovation = 5, total number of skills = 200. The number of skills possessed by agents is set as follows: LKB = 35, MKB = 45 and HKB = 55. Moreover, the rewiring parameter p was varied as values drawn from the set: {0.1, 0.3, 0.5}. Hence, as we did in Chapter 4, we specify three different simulations,⁸ whose results are presented hereafter.

At the outset, we want to look at the innovation patterns in the three simulations. In Figure 7.2 we report the innovation curves of both total and joint innovations. The system reaches the steady state in less than 15 time steps achieving just 11 innovations, out of the 120 available. This poor performance did not allow the network to evolve (no extra links were added as the system never reached 10 per cent of the innovations) and, therefore, the three simulations are identical. Recalling that the rewiring parameter was increasing




Figure 7.2 *Total and joint number of innovations*

the number of connections in the system as new innovations were attained, we can conclude that the poor performance of the system prevents the rewiring parameter from showing its positive effect. This preliminary finding is in line with the main stylized facts emerging from the empirical investigations conducted by Morone et al. (2006) and discussed in Chapter 5, where a poor performance in terms of knowledge-based exchanges was observed among Foggia's organic producers.

We should now try to understand the reasons behind such poor performance. As we observed in the theoretical model, two key variables affecting the system performance are the density of the acquaintances network and the size of its largest component. As both these variables measure the opportunity for integrating knowledge, they have direct implications on the innovative performance of the system.

In Table 7.2 we report a summary of the main network statistics for the acquaintances and the partnership network. As we can immediately observe, the acquaintances network is static over time, all network statistics being identical at the beginning and at the end of the simulation. This is a highly disconnected system with just 14 links connecting a population of 32 firms. It follows that the density is very low and the largest component is small in size. The partnership network is particularly small as it reaches a maximum of four

Table 7.2 Summary statistics for the acquaintances and partnership networks

	Number of edges	Number of links	Density	Cliquishness	Size of the largest component
<i>Acquaintance network</i>					
Beginning of simulation	32	14	0.014	0.375	7
End of simulation	32	14	0.014	0.375	7
<i>Partnership network</i>					
Beginning of simulation	32	2	0.002	0.125	2
End of simulation	32	4	0.004	0.125	4

links at the end of the simulation. This, of course, is reflected in a low density, low cliquishness and a small largest component.

As discussed in Chapter 4, it looks as though the system is locked in to an underperforming pathway which is determined by the low density of the system. In other words, the relatively small number of connections existing among the firms involved in the innovation process undermines their ability to perform well in terms of innovations. In order to test this hypothesis we will run new simulations, adding more links between firms. This will change the network configuration, increasing the system connectivity and its density. However, adding new links can reshape the network in rather different ways depending on where such new links are added. Hence, this time we will carry out repeated simulation experiments (batches of 100 runs each), to dispose of artefacts introduced by the random aspect of network reconfiguration. We will then look at the average performance over the batch.

In Table 7.3 we report the results obtained by increasing the number of links, specifying the number of innovations achieved at the end of the simulation both individually and jointly. We can immediately observe that, although there is an improvement as we increase the number of links, the system performance is still, overall, rather poor. Moreover, we can notice that our first simulation run

Table 7.3 System performance when increasing the network density (average results over 100 runs)

Number of edges	Number of links	Density	Innovations achieved at the end of the simulation					
			0.1		0.3		0.5	
			<i>II</i>	<i>JI</i>	<i>II</i>	<i>JI</i>	<i>II</i>	<i>JI</i>
32	14	0.014	3.74	3.56	3.74	3.56	3.74	3.56
32	24	0.024	4.05	6.34	3.99	6.34	3.94	6.28
32	34	0.034	4.17	7.25	4.13	7.43	4.34	7.78
32	44	0.044	5.19	9.35	5.33	9.95	5.47	10.11

is performing well when compared to the average performance of comparable runs (compare the first row of Table 7.3 with Figure 7.2).

All in all, these new simulations suggest that, under the prevailing condition of the organic producers operating in the province of Foggia, increasing the density of the system is not an effective policy action to increase their innovative performance. This outcome seems to counter earlier results obtained in Chapter 4, where we showed how the initial density of the system (along with the initial size of the largest component) was significantly affecting the final outcome of the simulation.

A possible explanation for such an outcome could be that firms are endowed with an insufficient knowledge base. In other words, the amount of skill present in the system is not sufficient to perform a large number of innovations even when the density of the system increases significantly.⁹ In order to test this hypothesis we shall increase the initial knowledge endowment of firms and compare their performances in terms of innovations achievements with earlier results. We increase the initial skill profile by adding ten extra skills to each firm. Hence we now have the following skills distribution: LKB = 45, MKB = 55 and HKB = 65.

In Table 7.4 we report the results obtained running the new simulation batches. As we can see, the system performs overall better;

Table 7.4 System performance when increasing the network density and the knowledge base (average results over 100 runs)

Number of edges	Number of links	Density	Innovations achieved at the end of the simulation					
			0.1		0.3		0.5	
			<i>II</i>	<i>JI</i>	<i>II</i>	<i>JI</i>	<i>II</i>	<i>JI</i>
32	14	0.014	4.69	8.26	4.69	8.26	4.93	8.48
32	24	0.024	4.88	9.46	5.01	9.65	4.93	9.47
32	34	0.034	6.46	13.69	7.26	14.51	6.98	15.36
32	44	0.044	10.52	20.95	10.22	21.71	10.92	22.43

however, a significant improvement arises only when adding 20 and 30 new links. Moreover, the system performance, although improved when compared to previous simulations, is still far from the upper bound of 120 new products.

As we keep increasing the initial knowledge base of firms we can notice that the system performance progresses in a non-linear fashion with the increase of the system’s density. As we can see in Figure 7.3, adding 20 extra skills produces a significant increase in the innovative performance, mainly when coupled with a denser network. Moreover, the performance differential observable when moving from less to more dense networks increases with the increase of the knowledge base up to 20 extra skills, at which point it starts reducing.

This finding suggests that the effectiveness of new connections between firms initially increases with the knowledge endowment of the cluster: increasing the network density exerts different effects, depending on the amount of knowledge present in the system, with more knowledgeable systems benefiting relatively more than less knowledgeable systems up to a certain level, and less subsequently. The trend becomes more evident as we increase the rewiring parameter p (contrast left and right panels in Figure 7.3).

Looking at Figure 7.3 we can also notice that comparable performances can be achieved by combining the two actions (that is,




Figure 7.3 Total number of innovations by various levels of initial knowledge and extra links


Table 7.5 Average number of innovations and standard deviation for two comparable outcomes obtained with different policy mixes ($p = 0.1$)

	Average	s.d.
Knowledge-oriented	52.59 (8.38)	60.78 (13.81)
Density-oriented	50.96 (18.15)	63.54 (18.46)

increasing firms’ knowledge base and increasing the density of the system) in different ways. For instance, with $p = 0.1$ a similar level of new products (around 50) can be achieved by adding 20 new skills and 20 new links as well as by adding only 30 new skills. Likewise, around 60 new innovations can be achieved either by adding 30 new skills and 10 new links or by adding 20 new skills and 30 new links. This would suggest that a policy-maker aiming at increasing the innovativeness of the system could achieve similar results with different mixes of the two policy actions.

However, we should recall that these are average results drawn from batches of 100 simulations each. Hence, in order to assess the effectiveness of such policy actions we should also look at the dispersion around the mean value. Interestingly, if we compare the standard deviation in all cases where the system achieves comparable results, we can notice that this tends to be larger whenever the mix of policy is more ‘density-oriented’ (that is, it involves a larger increase in the density of the network rather than an increase of its knowledge base). This suggests that a ‘knowledge-oriented’ policy action would produce more stable results, reducing the volatility of the outcome. In Table 7.5 we report the number of innovations, and the corresponding standard deviations, for the two exemplifying cases mentioned above.

This finding also suggests that a density-oriented action can produce exceptionally bad results as well as exceptionally good results (that is, further away from the average performance, see the grey line in Figure 7.4), this depending on which firms are connected. Hence, increasing the system density can produce better results if the right connections are activated. This is pointing at a nice feature of the model: establishing the right connections can be worth more



Note: Density-oriented batch: 20 extra links and 20 extra skills; knowledge-oriented batch: 0 extra links and 30 extra skills.

Figure 7.4 Single run performances around the average value, $p = 0.1$

than a generalized increase of firms' knowledge base, suggesting the importance of some specific linkages (or the effects of their absence) within innovation systems (and broader socio-economic systems) (Bryant, 2001). Hence, a well-crafted policy action, targeting the right firms and establishing the right connections, can be less costly and more effective than actions aiming at a generalized improvement of firms' knowledge base.

There are two possible approaches for assessing the validity of this last finding. First, we can define a general rule to activate new links in a more intelligent way rather than adding them randomly, run new batches of simulations, and inspect the average results. This will allow us to develop a broad-in-scope and robust assessment of the relevance of different rules for new links activations.

A second strategy entails an in-depth investigation of two comparable runs (that is, extracted from the same batch, hence obtained with the same parameterization but different initial network configuration) performing rather differently (that is, one run performing exceptionally well and one run performing exceptionally poorly). This will allow us to pinpoint specific structural differences between the good and the bad run and, in turn, to identify the key determinants of heterogeneous performances. In what follows we will pursue both approaches.

We start by defining three alternative ways of activating new links. We will run the same simulations as those depicted in Figure 7.3, adding new links: first departing from high knowledgeable firms, then departing from medium knowledgeable firms and, finally, departing from low knowledgeable firms. As we increase the number of connections moving from more to less knowledgeable firms, we expect to observe a reduction in the overall performance of the system. Looking at the results reported in Table 7.6 we can immediately see that in almost all cases the system performs better if the new connections start from knowledgeable agents. Hence, increasing the density of the system is more effective when the centrality of highly skilled firms increases. This is in line with our expectations and confirms our earlier finding that a more accurate selection of agents to be connected results in a better performance of the system.

We will now look at a good run and a bad run¹⁰ extracted from the same batch. Specifically we look at runs 65 (bad run) and 93 (good run) extracted from the batch with 20 new links, 20 new skills $p = 0.1$ and random selection of new links. First, we look at the

Table 7.6 Total number of innovations achieved using different rules to activate new links

	10 new links added				20 new links added				30 new links added			
	$p = 0.1$	$p = 0.3$	$p = 0.5$	$p = 0.1$	$p = 0.3$	$p = 0.5$	$p = 0.1$	$p = 0.3$	$p = 0.5$	$p = 0.1$	$p = 0.3$	$p = 0.5$
<i>Original KB</i>												
Random	10.39	10.33	10.22	11.42	11.56	12.12	14.54	15.28	15.58	14.54	15.28	15.58
LKB	10.19	10.19	10.61	11.56	11.45	11.90	14.96	14.55	15.42	14.96	14.55	15.42
MKB	9.77	9.72	9.52	11.65	11.74	12.11	15.83	15.73	15.55	15.83	15.73	15.55
HKB	10.37	10.33	10.47	14.68	15.39	15.63	20.23	20.43	21.65	20.23	20.43	21.65
<i>Original KB + 10</i>												
Random	14.34	14.66	14.40	20.15	21.77	22.34	31.47	31.93	33.35	31.47	31.93	33.35
LKB	11.42	11.98	11.97	14.37	14.75	14.76	19.60	20.48	21.12	19.60	20.48	21.12
MKB	13.36	13.99	14.42	21.79	22.50	23.95	31.67	31.30	37.09	31.67	31.30	37.09
HKB	17.68	17.85	17.93	33.15	30.59	34.76	47.21	47.87	56.09	47.21	47.87	56.09
<i>Original KB + 20</i>												
Random	37.21	38.43	40.74	50.96	55.10	58.04	63.54	63.09	71.82	63.54	63.09	71.82
LKB	31.54	32.06	34.63	37.84	38.10	39.38	42.65	46.10	49.72	42.65	46.10	49.72
MKB	39.35	40.81	42.19	55.58	60.14	63.72	69.53	73.22	77.71	69.53	73.22	77.71
HKB	43.12	43.16	46.55	56.65	58.51	59.37	69.21	75.44	76.66	69.21	75.44	76.66
<i>Original KB + 30</i>												
Random	60.78	63.34	64.52	69.87	74.01	77.69	78.42	82.66	84.21	78.42	82.66	84.21
LKB	57.87	59.27	63.81	62.38	67.20	69.25	67.82	71.55	74.82	67.82	71.55	74.82
MKB	58.40	62.11	65.05	68.04	72.62	74.22	75.97	77.79	81.39	75.97	77.79	81.39
HKB	68.76	71.16	75.09	81.78	84.81	86.75	91.60	91.46	94.60	91.60	91.46	94.60




Figure 7.5 Innovation performances (good run vs bad run)

innovation performance in the two runs. In Figure 7.5 we reproduce the total number of innovations and the number of joint innovations achieved in each run. As we can see, the good run takes about 40 time steps to converge towards the steady state, whereas the bad run converges to its stationary equilibrium in less than 20 time steps. Approximately 26 per cent of the total innovations achieved in the good run (that is, 25 out of 93) are obtained through individual innovation, whereas in the bad run this share does not exceed 20 per cent (that is, 4 out of 20).

Given that the initial knowledge base of the firms is the same in the two runs (that is, each firm is endowed with exactly the same skill profile in the two runs), this finding suggests that interactive learning, occurring while jointly innovating, is playing a crucial role in enhancing the ability of firms to innovate individually. In fact, if we look at the number of skills learned through interactions in the two runs, we can see that they add up to just 73 skills in the bad run, and 346 skills in the good run. This difference is due to the fact that interactive learning occurs any time there is a successful joint interaction (that is, any time two or more firms jointly innovate). As shown in Figure 7.5, the number of joint innovations in the good run is more than four times as large as the number of joint innovations (JI) in the bad run (that is, 16 JI in the bad run vs 68 JI in the good run). Hence, we can conclude that a good performance

Table 7.7 Summary statistics for the acquaintances and partnership networks (good run vs bad run)

Bad run	Number of edges	Number of links	Density	Cliquishness	Size of the largest component
<i>Acquaintance network</i>					
Beginning of simulation	32	34	0.034	0.307	27
End of simulation	32	34	0.034	0.307	27
<i>Partnership network</i>					
Beginning of simulation	32	1	0.001	0.063	2
End of simulation	32	11	0.011	0.406	3
Good run	Number of edges	Number of links	Density	Cliquishness	Size of the largest component
<i>Acquaintance network</i>					
Beginning of simulation	32	34	0.034	0.291	29
End of simulation	32	36	0.036	0.318	30
<i>Partnership network</i>					
Beginning of simulation	32	4	0.004	0.025	2
End of simulation	32	24	0.024	0.326	19

in terms of joint innovations exerts a positive externality on the capability of innovating individually, via interactive learning.

We shall now turn our attention to the determinants of joint innovations, attempting to identify the reasons behind the performance gap across the two runs. As we know, the two runs are identical, with the exception of the way in which the extra 20 links have been added. Hence, the explanation of the performance's differential must lie in the emerging different network structures.

In Table 7.7 we report a set of summary statistics for the two runs

registered at the beginning and at the end of each simulation. As expected, the acquaintance networks are highly comparable: they have similar levels of density and cliquishness, and a largest component of a comparable size. Most of the differences emerge when we look at the partnership networks, the good run displaying a much wider size as well as a much wider largest component at the end of the simulation.

This difference stems from the fact that more innovations are performed in the good run. However, it also suggests that in the good run a larger number of firms are involved in the innovation process: nearly 80 per cent of the firms operating in the good run are involved in at least one innovation process, whereas this percentage drops to 50 per cent in the bad run.

Further insights can be obtained by looking at the distribution of innovations across firms. In the good run there are eight firms (namely firms 10, 12, 14, 15, 18, 20, 22 and 31) involved in two-thirds of the innovation activities, suggesting that such agents play a central role in shaping the overall performance of the system. In fact, these firms have established very effective partnerships which allow them to achieve a large number of innovations. These effective partnerships are not present in the bad run and result in a drastic drop in the productivity of such firms. A case in point is provided by firm 18 (see Table 7.7 where we report the final configuration – that is, the configuration reached in the steady state equilibrium – of the partnership networks of both good runs and bad runs) which is the best-performing firm in the good run (along with firm 15), participating in 18 innovations, but does not take part in any innovation in the bad run.

This example points out that, in our model, connections are the real driver of innovations: in the good run firm 18 is connected, among others, to firms 22 and 15 (whereas these two links are both missing in the bad run), and successfully innovates six times with firm 22 and eight times with firm 15; the remaining four innovations are achieved individually thanks to knowledge spilled over from previous interactions with these two partners. Other examples of successful partnerships, present in the good run but not in the bad run, are those established between firm 14 and firm 12 and between firm 15 and firm 30 (see Figure 7.6) which resulted in a total of seven new products.

Finally, it is interesting to observe how the initial knowledge base of firms is not affecting their innovative performance (see Figure

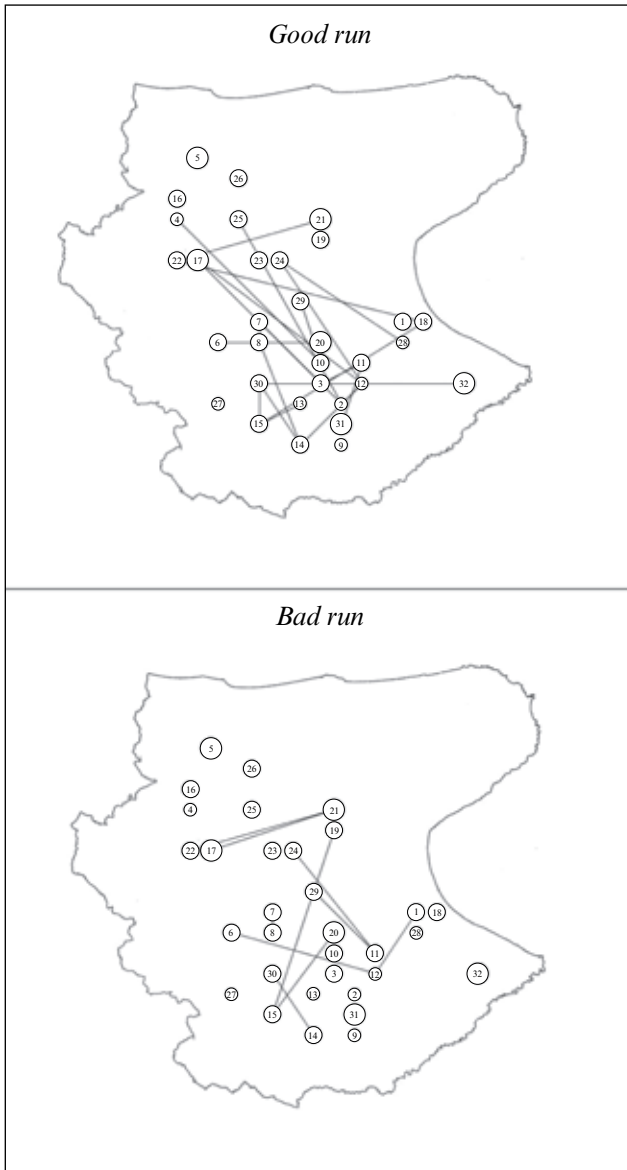


Figure 7.6 Partnership network graph, steady state configuration

7.6). For instance, firm 12 has a low KB but yet participates in 15 innovations. This confirms, once again, that links (or their absence) is what really shapes firms' innovative performances.

CONCLUSIONS

In this chapter we have presented a calibrated agent-based model of knowledge diffusion and innovation. This is an applied version of the theoretical model presented and discussed in Chapter 4 which offered us the opportunity to test the adherence of our theoretical model to an applied case study.

Overall, the results of the applied model confirmed the general trends observed in the theoretical model and the main stylized facts emerging from the case study performed in Morone et al. (2006) and discussed in Chapter 5. However, the calibrated model presented in this chapter also allowed the pinpointing of some features of the system specific to the case study. In particular, we could assess the impact of different policy actions aiming at enhancing the innovative performance of the local system. Results of the applied agent-based model showed that establishing the right connections was more effective than a generalized increase of firms' knowledge base, and pointed out the importance of linkages (or the effects of their absence) within innovation systems.

In fact, the presence of few highly effective partnerships exerted several positive effects on the firms' cluster as a whole, increasing the amount of knowledge diffused informally (interactive learning) in the system and favouring the involvement of a larger number of firms in innovation activities. These, in turn, resulted in an enhanced overall innovative performance. As far as policy-making is concerned, we could conclude that well-crafted policy action, targeting the right firms and enhancing the overall system connectivity, could be more effective and less costly than implementing generic training programmes extended to all firms.

NOTES

1. Strictly speaking, in the calibration exercise presented in this chapter we are not following all four pragmatic steps of indirect calibration discussed in Chapter 6.

However, we refer to it as indirect calibration since from the analysis of stylized facts performed in a previous work (Morone et al. 2006, discussed in Chapter 5), we believe that the model architecture developed in Chapter 4 is well suited to investigate innovation activities occurring in the system under empirical investigation. Moreover, as required in the indirect calibration approach, the parameters' space is restricted in accordance with the empirical data gathered in the case study.

2. Note that the relative populations of each of the 34 *comuna* composing the grid were respected in the calibration process.
3. For a detailed description on how to create such network structures please refer to Watts and Strogatz (1998).
4. Note that this cluster of firms is smaller than the one used in Morone et al. (2006) and discussed in Chapter 5, as it excludes those firms added following the free recall approach.
5. As stated in Chapter 5, the relevance of extra-cluster knowledge relations rests on the fact that the mere reliance on localized knowledge can result in the 'entropic death' of the cluster that remains locked in to an increasingly obsolete technological trajectory.
6. Note that the knowledge base (KB) of the firm is proxied by the amount of training provided to the employees of the firm. Specifically, we distinguished among four types of training (training course, participation in seminars and conferences, guided tours relevant to the production activity, other training activities) and classified as low KB those firms providing none, as medium KB those firms providing one or two types of training activities, and as high KB those firms providing three of four training activities.
7. Note that in this model we use a predetermined initial network which reflects the actual geographical locations of firms and the informal networks of entrepreneurs; whereas in the theoretical model presented in Chapter 4 firms were randomly allocated in a space (the grid) that represented acquaintance proximity (Fioretti 2001).
8. Note that, differently from what we did in Chapter 4, we do not run batches of simulations as we now have a fixed network structure. Hence the following results refers to single runs. Subsequently, we will randomly place additional links to increase the system density.
9. This would also explain why the simulations presented in Chapter 4 performed, on average, considerably better when compared to those presented in this chapter. In fact, although the average amount of skills endowment assigned to firms in the two models is comparable, the two systems differ both in terms of knowledge distribution across firms (that is, in Chapter 4 it follows a normal distribution, whereas in this model we defined just three levels of knowledge base – low, medium and high) and total skills endowment (that is, in the model presented in this chapter we have a smaller population of firms when compared to the population of firms defined in Chapter 4).
10. Following a procedure similar to the one used in Chapter 4, we select the good and bad runs from within a batch in the following way: first, we order all 100 runs in a list according to their average level of performance; then the good run is randomly selected from the top 10 per cent of the distribution (that is, among the ten best runs) and the bad run is randomly selected from the bottom 10 per cent of the distribution (that is, among the ten worst runs).

8. Final remarks and future research

WRAPPING UP IDEAS

In modern societies knowledge is rightly considered a key resource in promoting innovation and economic development. As was emphatically stated in a World Bank report a decade ago: ‘Knowledge is like light. Weightless and intangible, it can easily travel the world, enlightening the lives of people everywhere’ (World Bank 1999, p. 1). However, knowledge sharing is not as simple and straightforward as it would seem at first sight. This is because valuable knowledge is hard to codify and, therefore, requires experience, personal contacts and direct interactions in order to be shared.

The classic distinction between tacit and codified knowledge is, in fact, a major issue in understanding knowledge diffusion patterns. If tacit knowledge corresponds to the portion of knowledge that each person possesses but cannot tell,¹ then transferring it represents a problem. The magnitude of this problem depends crucially on two issues: first, how relevant tacit knowledge is for innovation (as opposed to codified knowledge) and, second, whether it is possible to improve the codifiability of tacit knowledge. Both issues have attracted the attention of researchers without the emergence of a clear consensus. However, there is agreement on the idea that tacit and codified knowledge flow in rather different ways. This calls for a deeper understanding of knowledge structure when discussing diffusion processes.

Another problem associated with knowledge sharing arises from the growing specialization that technical knowledge has undergone over the last 200 years and, even more markedly, over the last quarter of a century. As put by Brusoni et al. (2001): ‘Since the Industrial Revolution, the production of useful knowledge, like the production of artifacts, has become increasingly specialized and professionalized, with the continuous emergence of new and useful

disciplines and subdisciplines' (2001, p. 597). Today's production requires the input of a wide range of specialized knowledge, and this has profound and contrasting implications for knowledge diffusion patterns. On the one hand, the growing number of disciplines for the design, development and manufacturing of new products amplifies the need to rely on knowledge developed externally to the firm and, therefore, promotes knowledge sharing to complement in-house research and development efforts. On the other hand, transferring specialized knowledge is not feasible as it would be extremely inefficient due to learning costs associated to specialized knowledge transfer. Hence, firms act to create links with other firms and/or institutions through which they can integrate disparate and specialized knowledge needed to innovate. This calls for a deep understanding of firms' partnership networks and, more generally, the geographical dimension of knowledge flows.

SUMMARY OF THE BOOK

Both these issues have been taken up in this book in order to define a unifying theory of knowledge diffusion. In Chapter 2 we presented a taxonomy of knowledge flows which stems directly from the tacit-codified distinction. In the proposed taxonomy we made a first distinction between knowledge gain and knowledge diffusion. The former relates to those processes of knowledge flows which deliberately involve a barter among subjects; the latter refers to unintended knowledge flows which can be economically exploited by the recipient agent. Knowledge gain does not require geographical proximity among actors since the exchanged knowledge is mainly codified and can be easily transferred (making use, for instance, of information and communication technologies).

Conversely, the kind of knowledge being diffused is typically tacit in nature, and requires direct interactions (as it is spatially sticky) as well as sufficient 'absorptive capacity' to be effectively recombined in the cognitive framework of the recipient agent. Knowledge gain was subsequently decomposed in knowledge exchange and trade, whereas knowledge spillover, transfer and integration were defined as subclasses of knowledge diffusion.

In Chapter 3, then, we took a further step towards the definition of a theory of knowledge diffusion, presenting a survey of some of the

most relevant theoretical studies on knowledge diffusion. We started by presenting some of the earliest diffusion models (that is, epidemic models) which were reconciled to knowledge diffusion under the double assumption that innovation needs information (on its existence and its value) to be diffused and that information equates to knowledge. Subsequently, we presented a series of more sophisticated diffusion models which, making use of game theory, accounted for heterogeneous agents with heterogeneous beliefs. However improved, such later studies suffered from a conceptual limitation which rests on the dichotomous definition of learning. The acknowledgment of such problems led us to introduce a new class of models which, from a different methodological perspective, study knowledge flows patterns. This class of model makes use of agent-based simulations and, basically, counters the logical argument for diffusion followed in earlier models: knowledge is now considered as an input of innovation which can be acquired by direct interactions. Hence, firms share knowledge to innovate and not, as assumed in earlier studies, knowledge on innovations. This is, in our view, a critical step forward as it shifts the attention from innovation to knowledge as a key input for innovation.

Moving along the path set by this second class of models, in the fourth chapter we presented an agent-based simulation model which provided an original attempt to establish the complex relations linking knowledge-sharing patterns, firms' partnering and the innovative capability of firms. Using an agent-based approach allowed us to capture the dynamics and complexity present in the diffusion model. The objective was to understand the dynamics which lead firms to partner together and jointly innovate. The unit of the analysis was the firm, whose knowledge domain was defined as an articulated and complex structure in which more specialized knowledge was pegged to less specialized knowledge. Such a structure, labelled skill profile, was allowed to branch and, therefore, firms could specialize in different fields. The learning process through which firms acquired their initial skills resembled their in-house research and development efforts.

Heterogeneous firms, endowed with different skills, would most likely need to integrate their knowledge in order to innovate. Innovation was defined as the vector of skills required to create a new product. If one firm possessed all required skills to accomplish an innovation, then this firm would innovate individually. However, if the firm was not able to innovate individually, it would try to

partner with its acquaintances and jointly innovate. Partnering happened through direct interactions among neighbour firms. Firms' neighbourhoods were created at the initialization phase when firms were randomly allocated over a grid and were allowed to establish a connection with those firms located within their visibility range (we used a Moore neighbourhood structure with a grid size equal to 20, visibility equal to two and a total population of 40 firms).

Once a partnership was established, the firm and the partner tried to achieve the selected innovation by integrating their respective skills. If such an integration process was not effective (that is, the innovation was not achieved), the search process would continue expanding the partnership to other firms. Note that we allowed firms to increase their number of connections over time, using a rewiring probability p (that is, the percentage of additional links connecting distant firms) which was initially set equal to 10 per cent. New links were added gradually, at the same pace as innovations were attained. Several batches of simulations were run (each batch was composed of 100 runs), reconfiguring each time the network structure and keeping initial parameterization constant.

Some interesting results were drawn from the simulation exercises. First, looking at the system performance we observed heterogeneous outcomes: the number of achieved innovations varied quite significantly across different runs. Investigating the causes of such an outcome led us to establish a positive relation between the initial density of the network and the steady state performance of the system. This finding suggested that geography had a role in shaping innovative patterns: networks where firms were located close to each other resulted in more innovative environments and outperformed sparse systems. This outcome was consistent with the earlier assumption that tacit knowledge is spatially sticky and requires proximity to flow among actors.

Subsequently, we investigated how this result varied, increasing the percentage of new links added when innovating. Letting p reach first 30 and then 50 per cent we observed that the gap between well-performing runs and badly performing runs grew larger. This came as a surprise, as one would expect that a sparse network would benefit more from the introduction of distant connections. However, the poor performance in terms of innovation held back the positive impact brought about by higher values of p and limited the network possibilities to increase its density. This, in turn, locked in the system

to an underperforming equilibrium. On the contrary a positive feedback effect occurred in the case of well-performing runs, enhancing their innovative performance. Looking at the network configuration we confirmed these findings, showing how increasing the rewiring probability allowed well-performing runs to establish a relatively high number of distant connections, as opposed to badly performing runs which maintained prevalently localized neighbourhood structures.

Chapter 4 concluded the theoretical part of this book, providing new insight to the issue of defining knowledge diffusion patterns and innovations. However, several questions arose from the discussion developed in the first part of the book. For instance, having defined knowledge as a complex phenomenon, how do we measure it? And more importantly, how could we measure informal and complex patterns of knowledge diffusion (such as knowledge integration, as depicted in Chapters 3 and 4)?

To such relevant questions we attempted answers in the second part of the book, which dealt with empirical and applied studies on knowledge diffusion. It began with a survey of recent empirical studies on knowledge diffusion. Departing from the consideration that the available data mainly refer to the notion of information and/or provide just too vague proxies of knowledge, several authors have attempted to measure knowledge flows in a more direct and reliable way. A major effort was made to produce an accurate definition of the type of social relations which lead firms to cooperate and share knowledge. Social network analysis provided researchers with a powerful tool to achieve such an aim. Using this tool made it possible to establish the role played by specific actors in sharing knowledge and allowed us to identify and measure inward and outward knowledge flows. Indeed, most of the recent research in this area has relied on case studies collecting accurate data 'in-the-field' using questionnaires and interviews.

Such in-depth studies have allowed us to underpin some key features of knowledge diffusion patterns, such as the role played by knowledgeable actors (represented by firms with a large knowledge base) in facilitating intra- and extra-cluster knowledge flows, the existence of various relational networks for various types of knowledge shared (which could span from technical, to juridical, to pure tacit, and so on), the mechanisms through which knowledge is integrated among colleagues operating in the same firm, and various other features.

Although these studies represent major advances in the empirical understanding of informal knowledge diffusion mechanisms, we observed how their adherence to the case study undermines the possibility of drawing general conclusions. This problem suggested the need to explore alternative approaches to the empirical investigation of knowledge diffusion.

We came back to this issue in Chapter 6, where we presented a methodological discussion on the core distinction between theoretical and applied models, drawing this distinction within the realm of agent-based simulation models. In this chapter we identified foundational modelling as that which is engaged with modelling a target area of social theory. Formalization of the theory involves clarification of the theory and development of conceptual models able to be implemented in a computer program. Experiments are carried out with the objective of gaining a better understanding of the theory by, for example, confirming or refusing an existing hypothesis.

Rather different is the approach followed by empirical modellers. Here the objective is to apply the methodology to a well-defined study where there is substantial access to field data. It is natural that the abstraction step involves reducing complexity in comparison with the real-world system, and by so doing developing an applied model that, by the nature of its being simpler and therefore more amenable to study, might produce some insight into the more complex real-world counterpart.

The discussion on theoretical and applied models was followed by a closely related analysis of validation of agent-based models. Validation was considered quite broadly, encompassing both inputs and outputs to the modelling as well as all stages of the model building and analysis. It related to both theoretical and applied models. The relevance of validating agent-based models stems from the fact that researchers need to be sure that model outcomes reflect persistent (locally generic) aspects of the system under study, rather than a modeller's choice of parameter settings, initial conditions, or software platform preferences. In spite of the relevance of this issue, computer-simulated models lack a common methodological protocol for model validation. Whilst it might not be possible to build up a definitive approach to validating agent-based models, we presented a classification of the main procedures that modellers can follow to establish the level of confidence that can be placed in the model findings and, therefore, improve that level of confidence placed in such models.

This methodological discussion paved the way for the applied model developed in the seventh chapter of the book, which was an applied version of the model presented in Chapter 4, validated by means of calibration. The calibration exercise involved the collection of primary data on a group of firms located in the south of Italy, operating in the food sector and involved in the production of organic food. The crafting of the model was identical to the theoretical model presented in Chapter 4 with the exception that in the calibrated model we had a network of 32 firms whose connections corresponded to the real links observed and recorded in the field work. We also used data on firms' knowledge bases in order to create the initial skill profile of each firm.

The applied model confirmed the general trends observed in the theoretical model, but allowed us to pinpoint some features of the system specific to the case study. In particular, we could assess the impact of different policy action aiming at enhancing the innovative performance of the local system. From this investigation we could conclude that establishing the right connections was worth more than a generalized increase of firms' knowledge bases, pointing out the importance of linkages (or the effects of their absence) within innovation systems. As a policy implication we could conclude that a well-crafted policy action, targeting the right firms and enhancing the overall system connectivity, could be more effective and less costly than implementing generic training programmes extended to all firms. Both learning and innovations could be achieved more effectively by means of direct interactions among producers.

FUTURE TRENDS AND FUTURE RESEARCH

This book attempts to bring new insights on the widely debated topic of the knowledge economy and, specifically, on the relation existing between knowledge diffusion patterns and innovative activities. In spite of the documented progresses made in recent years in understanding the full complexity of the problem, we feel this is an area of investigation which deserves further attention. Future research should, in our view, consider both theoretical studies and empirical investigations, as they are complementary to each other.

Theoretical investigations should, for instance, aim at clarifying further the processes of knowledge flows considering the role

of institutions. In general, interactions need institutions (such as markets). However, considering knowledge sharing, markets might fail to provide the right incentives and additional institutions will be needed (Steiner 2003). Hence, efforts should be directed to understand which institutions might provide the right incentives for collaboration and knowledge sharing. We indicate two areas of research where the role of institutions should be further investigated.

First, moving from the idea that connectivity and knowledge sharing cannot be effectively coordinated by conventional markets (Helmstädter 2003), some scholars have recently pointed at clusters as learning organizations which could effectively deploy non-market devices by which firms could coordinate their activities with other firms and other knowledge-generating institutions. Thus, clusters add up to more than an agglomeration of firms and should be regarded as informal institutions facilitating the cooperation between firms and public, semi-public and private research and development institutions (Steiner 2004). In this sense, the role of clusters should be further investigated and researchers should make extra efforts in shaping and modelling the institutional nature of such local environments.

Second, as just mentioned, public and private research institutions play a vital role in prompting knowledge creation and diffusion, and this is even more true in some knowledge-intensive sectors where universities play a major role. As pointed out in a recent paper by Rosell and Agrawal (2009), in just 14 years, from 1980 to 1993, the number of patents issued annually to US universities increased by 316 per cent, from 390 to 1622. Understanding how and how much of this knowledge produced outside the firm flows into private companies is a rather relevant issue, which has serious policy implications if we consider that in the policy arena, ultimate economic benefits are increasingly seen as the primary policy motivation for the public support of scientific research (Jaffe and Trajtenberg 1996). Bearing this in mind, knowledge diffusion models should be extended to integrate knowledge flowing from public and private research institutions into firms. This would add to the existing theoretical literature, introducing a new level of heterogeneity among actors since research institutions and private companies behave differently and pursue different goals.

The empirical literature has encountered several problems in addressing the issue of knowledge diffusion. As discussed broadly in Chapter 5, most of these problems stem from difficulties associ-

ated with measuring the flow of an intangible asset like knowledge. A large number of empirical studies on knowledge flows made use of patent data, using patent citations as a proxy for knowledge diffusion. However, these types of data do not capture the whole magnitude of the phenomenon as not all innovations are patentable, and neither are all patentable innovations chosen to be patented. To overcome such problems several researchers have performed in-depth case studies, using ethnographic methods to capture better the dynamics of knowledge flows and generate new insights into communication patterns.

However, we would also like to emphasize the contribution that modelling can provide to the study of knowledge diffusion. Whilst applied agent-based modelling presents researchers with a powerful tool which can be exceptionally useful in understanding complex dynamics, such models have high input information demands and are difficult to validate with conventional tactics. Conclusions and policy relevance tend to be quite specific to the case investigation. This suggests the need for more abstract models which can provide results more adherent to a general theory. Our personal feeling is that this simplicity is a desirable feature of the more theoretical models, as it facilitates a better understanding of the modelled processes.

One aspect that we have highlighted in the discussion of validation protocols is the different ways in which model implementations can be compared. In our view, future research needs to theorize better how a dual modelling approach can be carried out. How, in other words, theoretical and applied models can be compared in terms of model Data Generation Process (mDGP) and in terms of simulation outputs. The utility of developing a suite of related models in this manner should be further investigated, and from a technical point of view it would be useful to reflect on how difficult and time-intensive this is.

There is potential that further modelling work could support the analysis of the role of institutions, of geographic aspects of firms' interaction networks, of the interaction between intellectual property regimes and innovation systems, for example. Certainly, further efforts are required to develop such models and model methodologies, which represent an important challenge for future research.

NOTE

1. As discussed in Charter 2, in Polanyi's view, we know more than we can tell.

References

- Allen, R.C. (1983), 'Collective invention', *Journal of Economic Behaviour and Organization*, **4** (1), 1–24.
- Amin, A. and P. Cohendet (2004), *Architectures of Knowledge: Firms, Capabilities and Communities*, Oxford: Oxford University Press.
- Ancori, B., A. Bureth and P. Cohendet (2000), 'The economics of knowledge: the debate about codification and tacit knowledge', *Industrial and Corporate Change*, **9** (2), 255–87.
- Antonelli, C. (1996), 'Localized knowledge percolation processes and information networks', *Journal of Evolutionary Economics*, **6** (3), 281–95.
- Arora, A., A. Fosfuri and A. Gambardella (2002), *Markets for Technology: The Economics of Innovation and Corporate Strategy*, Cambridge, MA: MIT Press.
- Arthur, W.B., J. Holland, B. LeBaron, R. Palmer and P. Tayler (1997), 'Asset pricing under endogenous expectations in an artificial stock market', in W.B. Arthur, S. Durlauf and D. Lane (eds), *The Economy as an Evolving Complex System II*, Reading, MA: Addison-Wesley, pp. 15–44.
- Arrow, K.J. (1962), 'Economic welfare and the allocation of resources to invention', in R. Nelson (ed.), *The Rate and Dimension of Innovation Activity*, Princeton, NJ: Princeton University Press, pp. 609–26.
- Asheim, B.T. (1996), 'Industrial districts as "learning regions": a condition for prosperity', *European Planning Studies*, **4** (4), 379–400.
- Avermaete, T. and J. Viaene (2002), 'On innovation and meeting regulation – the case of the Belgian food industry', DRUID Summer Conference on *Industrial Dynamics of the New and Old Economy – Who is Embracing Whom?*, 6–8 June, Elsinore, Denmark.
- Axelrod, R. (1984), *The Evolution of Cooperation*, New York: Basic Books.

- Axtell, R., R. Axelrod, J.M. Epstein and M.D. Cohen (1996), 'Aligning simulation models: a case study and results', *Computational and Mathematical Organization Theory*, **1** (2), 123–41.
- Bala, V. and S. Goyal (1995), 'A theory of learning with heterogeneous agents', *International Economic Review*, **36** (2), 303–23.
- Bala, V. and S. Goyal (1998), 'Learning from neighbours', *Review of Economic Studies*, **65**, 595–621.
- Becattini, G. and E. Rullani (1993), 'Sistema locale e mercato globale', *Economia e Politica Industriale*, **80**, 25–49.
- Berends, J.J., H. van der Bij, K. Debackere and M. Weggeman (2006), 'Knowledge sharing mechanisms in industrial research', *R&D Management*, **36** (1), 85–95.
- Bernhardt, I. and K.M. MacKenzie (1972), 'Some problems in using diffusion models for new products', *Management Science*, **19** (2), 187–200.
- Boudouropoulos, I.D. and I.S. Arvanitoyannis (2000), 'Potential and perspectives for application of environmental management system (EMS) and ISO 14000 to food industries', *Food Reviews International*, **16** (2), 177–237.
- Breschi, S. and F. Lissoni (2001), 'Knowledge spillovers and local innovation systems: a critical survey', *Industrial and Corporate Change*, **10** (4), 975–1005.
- Breschi, S. and F. Lissoni (2003), 'Mobility and social networks: localised knowledge spillovers revisited', CESPRI Working Papers, WP no. 142.
- Brökel, T. and M. Binder (2007), 'The regional dimension of knowledge transfers: a behavioral approach', *Industry and Innovation*, **14** (2), 151–75.
- Brusoni, S., A. Prencipe and K. Pavitt (2001), 'Knowledge specialization, organizational coupling, and the boundaries of the firm: why do firms know more than they make?' *Administrative Science Quarterly*, **46** (4), 597–621.
- Bryant, K. (2001), 'Promoting innovation: an overview of the application of evolutionary economics and systems approaches to policy issues', in J. Foster and J. St Metcalfe (eds), *Frontiers of Evolutionary Economics, Competition, Self Organization and Innovation Policy*, Cheltenham, UK and Northampton, MA, USA: Edward Elgar, pp. 361–83.
- Burns, T. and G. Stalker (1961), *The Management of Innovation*, London: Tavistock.

- Cabrera, A. and E.F. Cabrera (2002), 'Knowledge-sharing dilemmas', *Organization Studies*, **23**, 687–710.
- Cairncross, F. (1997), *The Death of Distance*, Boston, MA: Harvard Business School Press.
- Camagni, R. (1991), 'Local milieu, uncertainty and innovation networks: towards a new dynamic theory of economic space', in R. Camagni (ed.), *Innovation Networks: Spatial Perspectives*, London: Belhaven Press, pp. 121–42.
- Cantner, U. and H. Graf (2006), 'The network of innovators in Jena: an application of social network analysis', *Research Policy*, **35** (4), 463–80.
- Cantwell, J. and S. Iammarino (2003), *Multinational Corporations and European Regional Systems of Innovation*, London, UK and New York, USA: Routledge.
- Casetti, E. and R. Semple (1969), 'Concerning the testing of spatial diffusion hypotheses', *Geographical Analysis*, **1**, 254–9.
- Cassi, L. and L. Zirulia (2008), 'The opportunity cost of social relations: on the effectiveness of small worlds', *Journal of Evolutionary Economics*, **18** (1), 77–101.
- Casti, J.L. (1997), *Would-Be Worlds: How Simulation Is Changing the Frontiers of Science*, New York: John Wiley & Sons.
- CAVES (2007), 'Working paper on case study structure, stakeholder/agents and validation data', Deliverable No. 8 of Project 012816: CAVES – *Complexity, Agents, Volatility, Evidence and Scale*.
- Chwe, M.S.-Y. (2000), 'Communication and coordination in Social networks', *Review of Economic Studies*, **67**, 1–16.
- Cimoli, M. and M. Della Giusta (2000), 'The nature of technological change and its main implications on national and local systems of innovation', in P. Batten and D. Martellato (eds), *Innovation and Regional Development*, Boston, MA, USA; Dordrecht, Netherlands and London, UK: Kluwer Academic.
- Collins, H. (2001), 'Tacit knowledge, trust and the Q of sapphire', *Social Studies of Science*, **31** (1), 71–85.
- Cooke, P. (2001), *Knowledge Economies: Clusters, Learning and Co-Operative Advantage*, London: Routledge.
- Cooke, P. and K. Morgan (2000), *The Associational Economy: Firms, Regions and Innovation*, Oxford: Oxford University Press.
- Cowan, R. (2004), 'Network models of innovation and knowledge diffusion', MERIT-Infonomics Research Memorandum Series, N.2004–016.

- Cowan, R., P.A. David and D. Foray (2000), 'The explicit economics of knowledge codification and tacitness', *Industrial and Corporate Change*, **9**, 211–53.
- Cowan, R. and D. Foray (1997), 'The economics of codification and the diffusion of knowledge', *Industrial and Corporate Change*, **6**, 595–622.
- Cowan, R. and N. Jonard (2003), 'The dynamics of collective invention', *Journal of Economic Behavior and Organization*, **52** (4), 513–32.
- Cowan, R. and N. Jonard (2004), 'Network structure and the diffusion of knowledge', *Journal of Economic Dynamics and Control*, **28**, 1557–75.
- Crespi, G., C. Criscuolo, J.E. Haskel and M. Slaughter (2007), 'Productivity growth, knowledge flows and spillovers', CEP Discussion Paper No 785.
- David, P.A. and D. Foray (2002), 'An introduction to the economy of the knowledge society', *International Social Science Journal*, **54** (171), 9–23.
- Demsetz, H. (1991), 'The theory of the firm revisited', in O.E. Williamson, S.G. Winter and R.H. Coase (eds), *The Nature of the Firm*, New York: Oxford University Press, pp. 159–78.
- Dicken, P. and A. Malmberg (2001), 'Firms in territories: a relational perspective', *Economic Geography*, **77** (4), 345–63.
- Dosi, G. and L. Marengo (1994), 'Some elements of an evolutionary theory of organisational competences', in R. England (ed.), *Evolutionary Concepts in Contemporary Economics*, Ann Arbor, MI: University of Michigan Press, pp. 157–78.
- Dosi, G., O. Marsili, L. Orsenigo and R. Salvatore (1995), 'Technological regimes, selection and market structures', *Small Business Economics*, **7** (6), 411–36.
- Edmonds, B. and D. Hales (2003), 'Replication, replication and replication: Some hard lessons from model alignment', *Journal of Artificial Societies and Social Simulation*, **6** (4), available online at: <http://jasss.soc.surrey.ac.uk/6/4/11.html>.
- Edmonds, B. and S. Moss (2005), 'From KISS to KIDS: an "anti-simplistic" modelling approach', in P. Davidsson, B. Logan and T. Takadema (eds), *Multi Agent Based Simulation 2004*, Berlin: Springer, pp.130–44.
- Ellison G. and D. Fudenberg (1993), 'Rules of thumb of social learning', *Journal of Political Economy*, **101** (4), 612–43.

- Ellison G. and D. Fudenberg (1995), 'Word-of-Mouth Communication and Social Learning', *Quarterly Journal of Economics*, **110** (1), 93–125.
- Epstein, J.M. and R. Axtell (1996), *Growing Artificial Societies: Social Science from the Bottom Up*, Washington, DC: Brookings and Cambridge, MA: MIT Press.
- European Commission (1999), 'Evaluating socio-economic programmes: principal evaluation techniques and tools', *MEANS Collection, Volume 3*, Luxembourg: Office for Official Publication of the European Communities.
- Fagiolo, G., A. Moneta and P. Windrum (2007), 'Confronting agent-based models with data: methodological issues and open problems', in C. Bruun (ed.), *Advances in Artificial Economics*, Berlin: Springer, pp. 255–68.
- Fioretti, G. (2001), 'Information structure and behaviour of a textile industrial district', *Journal of Artificial Societies and Social Simulation*, **4** (4), available online at: <http://www.soc.surrey.ac.uk/JASSS/4/4/1.html>.
- Florida, R.L. (1995), 'Toward the learning region', *Futures*, **27** (5), 527–36.
- Foray, D. (2004), *Economics of Knowledge*, Cambridge, MA: MIT Press.
- Forero-Pineda, C. and H. Salazar (2002), 'The access of researchers from developing countries to international science and technology', *International Social Science Journal*, **54** (171), 129–40.
- FORFÁS (2004), *Innovation Networks*, Wilton Place, Dublin: Wilton Park House.
- Forsman, M. and N. Solidanter (2003), 'Knowledge transfer in clusters and networks', *Journal of International Business Studies – Literature Review*, **3**, available online at: http://copenhagen.jibs.net/LitReview/2003/2003_3_24.pdf.
- Fujita, M., P. Krugman and A.J. Venables (1999), *The Spatial Economy: Cities, Regions and International Trade*, Cambridge, MA: MIT Press.
- Geroski, P.A. (2000), 'Models of technology diffusion', *Research Policy*, **29**, 603–25.
- Gertler, M.S. (2003), 'Tacit knowledge and the economic geography of context, or the indefinable tacitness of being (there)', *Journal of Economic Geography*, **3** (1), 75–99.
- Gilbert, N., A. Pyka and P. Ahrweiler (2001), 'Innovation networks:

- a simulation approach', *Journal of Artificial Societies and Social Simulation*, **4** (3), available online at: <http://www.soc.surrey.ac.uk/JASSS/4/3/8.html>.
- Giuliani, E. and M. Bell (2005), 'The micro-determinants of meso-level learning and innovation: evidence from a Chilean wine cluster', *Research Policy*, **34** (1), 47–68.
- Grabher, G. (1993), 'The weakness of strong ties: the lock-in of regional development in the Ruhr area', in G. Grabher (ed.), *The Embedded Firm*, London: Routledge, pp. 1–32.
- Grant, R.M. (1996), 'Prospering in dynamically-competitive environments: organizational capability as knowledge integration', *Organization Science*, **7**, 375–87.
- Griliches, Z. (1957), 'Hybrid corn: an exploration in the economics of technological change', *Econometrica*, **25**, 501–22.
- Grunert, K.G., K. Brunsø and S. Bisp (1997), 'Food-related life style: development of a crossculturally valid instrument for market surveillance', in L.R. Kahle and L. Chiagouris (eds), *Values, Lifestyles and Psychographics*, Mahwah, NJ: Lawrence Erlbaum, pp. 337–54.
- Guerrieri, P., S. Iammarino and C. Pietrobelli (2001), *The Global Challenge to Industrial Districts*, Cheltenham, UK and Northampton, MA, USA: Edward Elgar.
- Hagedoorn, J. (2002), 'Inter-firm partnerships: an overview of major trends and patterns since 1960', *Research Policy*, **31**, 477–92.
- Hägerstrand, T. (1967), *Innovation Diffusion as a Spatial Process*, Chicago, IL: University of Chicago Press.
- Haldin-Herrgard, T. (2000), 'Difficulties in diffusion of tacit knowledge in organisations', *Journal of Intellectual Capital*, **1** (4), 357–65.
- 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.
- Helmstädter, E. (ed.) (2003), *The Economics of Knowledge Sharing, A New Institutional Approach*, Cheltenham, UK and Northampton, MA, USA: Edward Elgar.
- Howells, J.R.L. (2002), 'Tacit knowledge, innovation and economic geography', *Urban Studies*, **39** (5–6), 871–84.
- Jaffe, A.B. (1986), 'Technological opportunity and spillovers of R&D: evidence from firms patents, profits, and market value', *American Economic Review*, **76**, 984–1001.
- Jaffe, A.B., M. Trajtenberg and R. Henderson (1993), 'Geographical

- localization of knowledge spillovers as evidenced by patent citations', *Quarterly Journal of Economics*, **108**, 577–98.
- Jaffe, A.B. and M. Trajtenberg (1996), 'Flows of knowledge from universities and federal laboratories: modeling the flow of patent citations over time and across institutional and geographic boundaries', Working Paper No. 5712.
- Java Universal Network/Graph (JUNG) Framework, available at: <http://jung.sourceforge.net>.
- Johnson, B.H. and B.-Å. Lundvall (1994), 'The learning economy', *Journal of Industry Studies*, **1** (2), 23–42.
- Johnson, B.H., B.-Å. Lundvall and E. Lorenz (2002), 'Why all this fuss about codified and tacit knowledge?' *Industrial and Corporate Change*, **11** (2), 245–62.
- Kiefer, N. (1989), 'A value function arising in economics of information', *Journal of Economic Dynamics and Control*, **7**, 331–7.
- King, G. and L. Zeng (2001), 'Logistic regression in rare events data', *Political Analysis*, **9**, 137–63.
- Kleijnen, J.P.C. (1999), 'Validation of models: statistical techniques and data availability', in P.A. Farrington, H.B. Nembhard, D.T. Sturrock and G.W. Evans (eds), *Proceedings of the 1999 Winter Simulation Conference*, San Diego, CA: Society for Computer Simulation, pp. 647–54.
- Krugman, P. (1996), *Self-Organizing Economy*, Cambridge, MA: Blackwell.
- Laffont, J.J. (1989), *The Economics of Uncertainty and Information*, Cambridge, MA: MIT Press.
- Leombruni, R., M. Richiardi, N.J. Saam and M. Sonnessa (2006), 'A common protocol for agent-based social simulation', *Journal of Artificial Societies and Social Simulation*, **9** (1), available on-line at: <http://jasss.soc.surrey.ac.uk/9/1/15.html>.
- Lundvall, B.-Å. (1988), 'Innovation as an interactive process: from user–producer interaction to national systems of innovation', in G. Dosi, C. Freeman, R.R. Nelson, G. Silverberg and L. Soete (eds), *Technology and Economic Theory*, London: Pinter Publishers, pp. 349–69.
- Lundvall, B.-Å. and D. Foray (1998), 'The knowledge-based economy: from the economics of knowledge to the learning economy', in D. Neef, G.A. Siesfeld and J. Cefola (eds), *The Economic Impact of Knowledge*, Woburn, MA: Butterworth-Heinemann, pp. 115–22.

- Lundvall, B.-Å. and P. Maskell (2000), 'Nation states and economic development: from national systems of production to national systems of knowledge creation and learning', in G.L. Clark, M. Feldman and M. Gertler (eds), *The Oxford Handbook of Economic Geography*, Oxford: Oxford University Press, pp. 353–72.
- Mahajan, V. and R.A. Peterson (1985), *Models for Innovation Diffusion*, Beverly Hills, CA: Sage.
- Malerba, F., R.R. Nelson, L. Orsenigo and S.G. Winter (1999), 'History friendly models of industry evolution: the computer industry', *Industrial and Corporate Change*, **8**, 3–41.
- Malerba, F. and L. Orsenigo (2001), 'Innovation and market structure in the dynamics of the pharmaceutical industry and biotechnology: towards a history friendly model', *Conference in Honour of Richard Nelson and Sydney Winter*, Aalborg, Denmark, 12–15 June.
- Malmberg, A. (1997), 'Industrial geography: location and learning', *Progress in Human Geography*, **21** (4), 573–82.
- Malmberg, A. and P. Maskell (2002), 'The elusive concept of localization economies: towards a knowledge-based theory of spatial clustering', *Environment and Planning A*, **34** (3), 429–49.
- Marks, R.E. (2007), 'Validating simulation models: a general framework and four applied examples', *Computational Economics*, **30** (3), 265–90.
- Marshall, A. (1920), *Principles of Economics*, 8th edition, London: Macmillan.
- Maskell, P. and A. Malmberg (1999), 'The competitiveness of firms and regions: "ubiquitification" and the importance of localised learning', *European Urban and Regional Studies*, **6**, 9–25.
- Meder, A. (2008), 'Technological and geographical patterns in the choice of cooperation partner', Jena Economic Research Papers in Economics 2008-054, Friedrich-Schiller-University Jena, Max-Planck-Institute of Economics, Thueringer Universitaets- und Landesbibliothek.
- Morgan, K. (1997), 'The learning region: institutions, innovation and regional renewal', *Regional Studies*, **31**, 491–504.
- Morgan, K. (2004), 'The exaggerated death of geography: learning, proximity and territorial innovation systems', *Journal of Economic Geography*, **4**, 3–21.
- Morone, A., P. Morone and R. Taylor (2007), 'A laboratory experiment of knowledge diffusion dynamics', in U. Cantner and F.

- Malerba (eds), *Innovation, Industrial Dynamics and Structural Transformation*, Berlin: Springer, pp. 283–302.
- Morone, P., R. Sisto and R. Taylor (2006), ‘Knowledge diffusion and networking in the organic production sector: a case study’, *EuroChoices*, **5** (3), 40–46.
- Morone, P. and R. Taylor (2004a), ‘Knowledge diffusion dynamics of face-to-face interactions’, *Journal of Evolutionary Economics*, **14**, 327–51.
- Morone, P. and R. Taylor (2004b), ‘Small world dynamics and the process of knowledge diffusion: the case of the metropolitan area of Greater Santiago De Chile’, *Journal of Artificial Societies and Social Simulation*, **7** (2), available online at: <http://jasss.soc.surrey.ac.uk/7/2/5.html>.
- Morone, P. and R. Taylor (2008), ‘Knowledge architecture and knowledge flows’, *Encyclopaedia of Information Science and Technology*, 2nd edition, Hershey, PA: IGI Global Publisher.
- Morris, S. (2000), ‘Contagion’, *Review of Economic Studies*, **67**, 57–68.
- Moss, S. (2001), ‘Intuition and observation in the design of multi agent systems’, CPM Report No. 01-84, Manchester Metropolitan University, UK.
- Nelson, R.R. (1959), ‘The simple economics of basic research’, *Journal of Political Economy*, **67**, 297–306.
- Nelson, R.R. and E. Phelps (1966), ‘Investment in humans, technological diffusion and economic growth’, *American Economic Review*, **56** (1–2), 69–75.
- Nelson, R.R. and N. Rosenberg (1993), ‘Technical innovation and national systems’, in R.R. Nelson (ed.), *National Innovation Systems: A Comparative Analysis*, Oxford: Oxford University Press, pp. 3–22.
- Nelson, R.R. and S. Winter (1982), *An Evolutionary Theory of Economic Change*, Cambridge, MA: Harvard University Press.
- Nonaka, I. and H. Takeuchi (1995), *The Knowledge-Creating Company*, Oxford: Oxford University Press.
- Nooteboom, B. (1999), ‘Innovation, learning and industrial organization’, *Cambridge Journal of Economics*, **23**, 127–50.
- North, M.J., N.T. Collier and J.R. Vos (2006), ‘Experiences creating three implementations of the RePast agent modeling toolkit’, *ACM Transactions on Modeling and Computer Simulation*, **16** (1), January, 1–25.

- O'Brien, R. (1992), *Global Financial Integration: The End of Geography*, London: Royal Institute of International Affairs.
- O'Madadhain, J., D. Fisher, P. Smyth, S. White and Y. Boey (2005), 'Analysis and visualization of network data using JUNG', *Journal of Statistical Software*, **V** (1), 1–35.
- Oerlemans, L., M. Meeus and F. Boekema (1998), 'Learning, innovation and proximity', Eindhoven Centre for Innovation Studies, The Netherlands, Working Paper 98.3.
- Oreskes, N., K. Shrader-Frechette and K. Belitz (1994), 'Verification, validation, and confirmation of numerical models in the earth sciences', *Science*, **263**, 641–6.
- Pajares, J., A. Lopez and C. Hernandez (2003), 'Industry as organisation of agents: innovation and R&D management', *Journal of Artificial Societies and Social Simulation*, **6** (2), available online at: <http://www.soc.surrey.ac.uk/JASSS/6/2/7.html>.
- Pavitt, K. and P. Patel (1991), 'Large firms in the production of the world's technology: an important case of non-globalisation', *Journal of International Business Studies*, **22**, 1–21.
- Pinch, S., N. Henry, M. Jenkins and S. Tallman (2003), 'From "industrial districts" to "knowledge clusters": a model of knowledge dissemination and competitive advantage in industrial agglomeration', *Journal of Economic Geography*, **3**, 373–88.
- Polanyi, M. (1958), *Personal Knowledge*, London: Routledge & Kegan Paul.
- Polanyi, M. (1967), *The Tacit Dimension*, London: Routledge.
- Porter, M.E. (1998), 'Clusters and the new economics of competition', *Harvard Business Review*, **76** (6), 77–90.
- Pyka, A. (1997), 'Informal networking', *Technovation*, **17**, 207–20.
- Pyka, A. (2002), 'Innovation networks in economics: from the incentive-based to the knowledge-based approaches', *European Journal of Innovation Management*, **5** (3), 152–63.
- Rogers, E. (2003), *Diffusion of Innovations*, 5th edn, New York: Free Press, first published 1962, Glencoe: Free Press.
- Rosell, C. and A. Agrawal (2009), 'Have university knowledge flows narrowed? Evidence from patent data', *Research Policy*, **38** (1), 1–13.
- Rosenberg, N. (1982), *Inside the Black Box: Technology and Economics*, Cambridge: Cambridge University Press.
- Sahal, D. (1981), *Patterns of Technological Innovation*, New York: Addison Wesley.

- Saviotti, P.P. (1998), 'On the dynamics of appropriability, of tacit and of codified knowledge', *Research Policy*, **26** (7–8), 843–56.
- Saviotti, P.P. (1999), 'Knowledge, information and organisational structures', in P.L. Robertson (ed.), *Authority and Control in Modern Industry*, London: Routledge, pp. 120–39.
- Schelling, T.C. (1978), *Micromotives and Macrobehavior*, New York, USA and London, UK: W.W. Norton & Company.
- Senker, J. (1993), 'The contribution of tacit knowledge to innovation', *AI & Society*, **7**, 208–24.
- Silverberg, G., G. Dosi and L. Orsenigo (1988), 'Innovation, diversity and diffusion: a self-organization model', *Economic Journal*, **98**, 1032–54.
- Slicher van Bath, B.H. (1963), *The Agrarian History of Europe, AD 500–1850*, London: Arnold.
- Steiner, M. (2004), 'Knowledge networks as institutions for technology management', in C. Karlsson, P. Flensburg and S. Åke Hörte (eds), *Knowledge Spillovers and Knowledge Management*, Cheltenham, UK and Northampton, MA, USA: Edward Elgar, pp. 373–95.
- Steinmueller, E.W. (2000), 'Will new information and communication technologies improve the "Codification" of knowledge?' *Industrial and Corporate Change*, **9** (2), 361–76.
- Steinmueller, E.W. (2002), 'Knowledge-based economies and information and communication technologies', *International Social Science Journal*, **54** (171), 141–53.
- Zsulanski, G. (1996), 'Exploring internal stickiness: impediments to the transfer of best practice within the firm', *Strategic Management Journal*, **17** (winter), 27–43.
- Terán, O. (2001), 'Emergent tendencies in multi-agent-based simulations using constraint-based methods to effect practical proofs over finite sets of simulation outcomes', doctoral thesis, Manchester Metropolitan University, Manchester, available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=292408.
- van der Bij, H., X.M Song and M. Weggeman (2003), 'An empirical investigation into the antecedents of knowledge dissemination at the strategic business unit level', *Journal of Product Innovation Management*, **20**, 163–79.
- von Hippel, E. (1987), 'Co-operation between rivals: informal know-how trading', *Research Policy*, **16**, 291–302.
- von Hippel, E. (1988), *The Sources of Innovation*, Oxford: Oxford University Press.

- Wasserman, S. and K. Faust (1994), *Social Network Analysis. Methods and Applications*, Cambridge, MA: Cambridge University Press.
- Watts, D. and S. Strogatz (1998), 'Collective dynamics of small-world networks', *Nature*, **393**, 440–42.
- Werker, C. and T. Brenner (2004), 'Empirical calibration of simulation models', Papers on Economics and Evolution No. 0410, Jena: Max Planck Institute for Research into Economic Systems.
- Windrum, P. (2007), 'Neo-Schumpeterian simulation models', in H. Hanusch and A. Pyka (eds), *The Elgar Companion to Neo-Schumpeterian Economics*, Cheltenham, UK and Northampton, MA, USA: Edward Elgar, pp. 405–39.
- Windrum, P., G. Fagiolo and A. Moneta (2007), 'Empirical validation of agent-based models: alternatives and prospects', *Journal of Artificial Societies and Social Simulation*, **10** (2), available online at: <http://jasss.soc.surrey.ac.uk/10/2/8.html>.
- World Bank (1999), *World Development Report 1998/1999*, New York: Oxford University Press.
- Zeigler, B.P. (1985), *Theory of Modelling and Simulation*, Malabar, FL: Krieger; first published 1976, New York: Wiley.

Index

- absorptive capacity 92–3, 152
- agent-based modelling 33, 47, 81–2, 153
 - bilateral bartering model 34–7
 - complex model of knowledge
 - diffusion and innovation
 - defining firms' skills universe 51–2
 - definition of firms' skill profile 53
 - firms and their social network 50–51
 - firms' innovation and
 - partnerships 54–5, 154
 - interactive learning process 55
 - joint innovation and network change 55
 - main simulation phase 54
 - radical innovations versus
 - incremental innovations 52–3
 - simulation experiment 56–80, 154
 - firms' interactions, knowledge and innovation 48–50
 - modelling knowledge exchange
 - as costs and benefits
 - comparison 37–9
 - modelling knowledge transfer as
 - face-to-face diffusion 40–44
 - modelling knowledge transfer as
 - localized broadcasting 39–40
 - types of problem addressed with 105–10
 - validation of models 110–11
 - alternative approaches 116–17
 - conventional approach 111–13
 - protocols for validation 113–16
- Agrawal, A. 158
- agriculture, diffusion of innovation
 - in 25
- Allen, R.C. 32
- Ancori, B. 4, 10
- applied modelling 105
- Avermaete, T. 131
- Bala, V. 30–31
- bartering model 34–7
- Bell, M. 86, 92, 93, 95
- Berends, J.J. 19, 87, 97
- bilateral bartering model 34–7
- Binder, M. 16
- Brenner, T. 115
- Breschi, S. 16
- broadcasting
 - diffusion models 25–7
 - modelling knowledge transfer as
 - localized broadcasting 39–40
- Brökel, T. 16
- Brusoni, S. 151
- Cantner, U. 87, 88, 89, 91
- Cassi, L. 37, 38, 39
- Chile 86, 126, 127
- Chwe, M.S.-Y. 31
- closed systems 105
- clusters 158
- codified knowledge 11–14, 16, 18, 151
- cognitive distance 42–3
- collective invention 32
- comparison of models 117–22
- conceptual modelling 107, 109
- confirmation 106
- contagion 31
- cost–benefit comparison, modelling
 - knowledge exchange as 37–9

- Cowan, R. 13, 18, 34, 35–6, 37, 38, 39, 40, 41
 crude knowledge 4
- David, P.A. 8
 diffusion of innovation/knowledge 3–5, 18, 19–21, 152–3, 158–9
 complex model of knowledge diffusion and innovation
 defining firms' skills universe 51–2
 definition of firms' skill profile 53
 firms and their social network 50–51
 firms' innovation and partnerships 54–5, 154
 interactive learning process 55
 joint innovation and network change 55
 main simulation phase 54
 radical innovations versus incremental innovations 52–3
 simulation experiment 56–80, 154
 measuring knowledge diffusion 86–7
 in industrial research 97–100
 within innovation networks 87–92
 within social networks 92–7
 reviewing diffusion models 24–5
 epidemic models 25–9
 game-theoretical models 29–32
 validation of models 125–30, 149
 case study 130–32
- distance
 knowledge and 14
 social 16–17
- docking 118
- Dosi, G. 16
- Edmonds, B. 117
- Ellison, G. 30
- empirical studies on knowledge flows 85, 100–101, 156
 measuring knowledge diffusion 86–7
 in industrial research 97–100
 within innovation networks 87–92
 within social networks 92–7
 enrichment of knowledge base 7–8
 epidemic models of diffusion 25–9
 equality 35
 evolutionary approaches 31, 37
- face-to-face diffusion 40–44
 factual information 11
- Fagiolo, G. 33
- flow patterns of knowledge 17, 21
 decomposing knowledge diffusion 19–21
 knowledge gain versus knowledge diffusion 18
- models 32–4
 bilateral bartering model 34–7
 modelling knowledge exchange as costs and benefits comparison 37–9
 modelling knowledge transfer as face-to-face diffusion 40–44
 modelling knowledge transfer as localized broadcasting 39–40
- food industry case study of validation of knowledge diffusion models 130–32
 results of simulation model 132–49
- Foray, D. 8, 9, 10, 13
- foundational modelling 104, 105, 107, 109–10, 120, 156
- Fudenberg, D. 30
- future trends and future research 157–9
- game-theoretical diffusion models 29–32
 general theory 106

- geography, knowledge and 14–17, 49
- Germany, measuring knowledge
diffusion within innovation
networks in 87–92
- Geroski, P.A. 25, 26–7
- Giuliani, E. 86, 92, 93, 95
- globalization 14–15
- Goyal, S. 30–31
- Graf, H. 87, 88, 89, 91
- Grant, R.M. 3, 19
- Haldin-Herrgard, T. 13, 15
- Howells, J.R.L. 10
- human intellectual capital 3
- ignorance trap 43
- incremental innovations 52–3
- incremental learning 32
- industrial research, measuring
knowledge diffusion within
97–100
- inequalities 35
- information, definition 10–11
- innovation 3–5, 15
definition 7
diffusion, *see* diffusion of
innovation/knowledge
joint innovation and network
change 55
radical innovations versus
incremental innovations
52–3
- integration of knowledge 19–20
- interactive learning process 55
- invention, collective 32
- Italy 87
case study of validation of
knowledge diffusion models
130–32
results of simulation model
132–49
measuring knowledge diffusion
within social networks 96–7
- Jaffe, A.B. 24
- Johnson, B.H. 8, 13
- Jonard, N. 18, 34, 35–6, 37, 38, 39,
40, 41
- know-how 32
- knowledge 151–2
definition 8–9
see also individual topics
- laboratory experiments 122
- language 14
- learning 3, 10
incremental 32
interactive learning process 55
social 30–31
- learning organizations 8, 158
- Lissoni, F. 16
- localization 24
- logic 107
- Lundvall, B.-Å. 8, 9
- Malerba, F. 116
- Marengo, L. 16
- Marks, R.E. 110, 111, 112
- Marshall, Alfred 49, 86
- measuring knowledge diffusion
86–7
in industrial research 97–100
within innovation networks
87–92
within social networks 92–7
- Meder, A. 90, 91
- modelling 21–2, 24, 32–4, 44–5, 104,
159
agent-based, *see* agent-based
modelling
diffusion models 24–5
epidemic models 25–9
game-theoretical models
29–32
foundational modelling 104, 105,
107, 109–10, 120, 156
theoretical and applied modelling
104–5
validation, *see* validation of
models
see also empirical studies on
knowledge flows

- Morgan, K. 17
 Morone, P. 17, 41, 86–7, 96, 126, 128, 136
 Morris, S. 31
 Moss, S. 117

 Nelson, R.R. 12, 37
 neoclassical economics 24, 33
 networks 154–5
 joint innovation and network change 55
 measuring knowledge diffusion within innovation networks 87–92
 within social networks 92–7
 simulation experiment 59, 67, 69, 75
 small world networks 36
 social networks 37–8, 50–51, 155
 measuring knowledge diffusion within 92–7
 Nooteboom, B. 43

 Oreskes, N. 106
 Orseniga, L. 116

 partnerships, innovation and 54–5, 154
 patents 24, 158
 measuring knowledge diffusion within innovation networks 87–92
 perception 14
 Polanyi, M. 12
 Porter, M.E. 14–15
 public goods 24

 radical innovations 52–3
 reciprocity 16
 reproduction of knowledge 10–11
 research and development (R&D) 3, 8
 research institutions 158
 Rogers, E. 25

 role-playing games 122
 Rosell, C. 158
 Rosenberg, N. 12

 Saviotti, P.P. 11, 12–13
 Senker, J. 12
 sensitivity analysis 122
 simulation experiment 56–80, 154
 skills
 defining firms' skills universe 51–2
 definition of firms' skill profile 53, 153
 Slicher van Bath, B.H. 25
 small world networks 36
 social learning 30–31
 social networks 37–8, 50–51, 155
 measuring knowledge diffusion within 92–7
 specialization 151–2
 spillovers 18, 19, 24
 stakeholder validation 122–3
 Steinmueller, E.W. 13
 Strogatz, S. 39

 tacit knowledge 3–4, 11–14, 15, 16, 17, 32, 151
 Taylor, R. 17, 41, 126
 technology 15–16
 diffusion, *see* diffusion of innovation/knowledge
 theorem proving 123
 theoretical modelling 104–5
 time, knowledge and 14
 transfer of knowledge 19
 trust 16

 United States of America 158

 validation of models 105–6, 123–4, 156, 157
 agent-based models 110–11
 alternative approaches 116–17
 conventional approach 111–13

- protocols for validation
 - 113–16
- by comparing models 117–22
- further methods 122–3
- knowledge diffusion models
 - 125–30, 149
 - case study 130–32
 - results of simulation model 132–49
- Viaeneon, J. 131
- von Hippel, E. 16, 32–3
- Watts, D. 39
- Werker, C. 115
- Windrum, P. 113
- Winter, S.G. 12, 37
- word-of-mouth communication
 - 30
 - word-of-mouth diffusion models 25–6, 27–9
- World Bank 151
- Zirulia, L. 37, 38, 39

