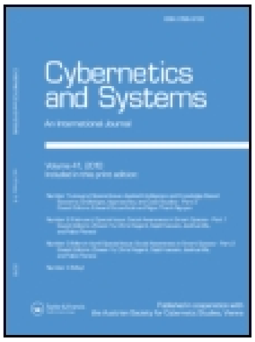

Cybernetics and Systems

An International Journal

ISSN: 0196-9722 (Print) 1087-6553 (Online) Journal homepage: https://www.tandfonline.com/loi/ucbs20

LEARNING-BY-MODELING: INSIGHTS FROM AN AGENT-BASED MODEL OF UNIVERSITY–INDUSTRY RELATIONSHIPS

Giorgio Triulzi & Andreas Pyka

Cybernetics and Systems: An International Journal, 42:484–501

ISSN: 0196-9722 print/1087-6553 online
DOI: 10.1080/01969722.2011.610266

Learning-by-modeling: Insights from an Agent-Based Model of University–Industry Relationships

GIORGIO TRIULZI[1] and ANDREAS PYKA[2]
[1]*UNU-MERIT, Maastricht University, Maastricht, The Netherlands*
[2]*University of Hohenheim, Stuttgart, Germany*

Learning is at the base of the so-called knowledge society. New forms of learning able to more effectively challenge the complex nature of phenomenon surrounding us are increasingly necessary. In this article we argue that learning-by-modeling, through a double-loop learning process, can significantly contribute to the refinement and improvement of our knowledge of complex phenomena. To sustain this argumentation, we make use of the insights provided by an agent-based model of university–industry relationships.

KEYWORDS agent-based modeling, double-loop learning, university–industry relationships

INTRODUCTION

In what is increasingly thought of as a "knowledge society" (World Bank 1999), continuous learning is claimed to be a fundamental requirement upon which knowledge and economic growth are built. In a complex world learning is not an easy task—on the contrary, learning is itself a complex phenomenon. In terms of a taxonomy, different forms of learning have been identified as so-called learning-by-concepts (Malerba 1992), of which so-called learning-by-doing is most prominent. A particular form of learning-by-doing, which concerns research activities, can be labeled *learning-by-modeling*: A problem of interest is described in an abstract form in an equation system and the static as well as dynamic features of the model

Address correspondence to Giorgio Triulzi, UNU-MERIT, Maastricht University, Keizer Karelplein 19, 6211 TC Maastricht, The Netherlands. E-mail: triulzi@merit.unu.edu

are analyzed. The insights gained are applied to generate a better understanding of the underlying problem.

Choosing an adequate mathematical representation as well as an adequate degree of complexity reduction to trace the model analytically is anything but an easy task. For instance, researchers in economics face a severe trade-off between the complexity of the problem under investigation, which is characterized by multi-agent dynamics and often self-reinforcing processes and the scope of an analytical framework that, at best, can give a detailed partial picture of a multifaceted reality, in particular concerning the dynamic features. In such contexts, agent-based modeling (ABM) is increasingly considered as a promising alternative (Tesfatsion and Judd 2006; Pyka and Fagiolo 2007); on the one hand, it builds upon theories and empirical evidence; on the other hand, it integrates them, suggesting different and more comprehensive ways to look at complex phenomena, thus providing the potential to explain results that otherwise might look contradictory. Modeling experience can be viewed as a learning process itself: it is an adaptive struggle in a world full of complexity. In this struggle we often become aware of how little we know about the actual problem under investigation.

The aim of this article is to discuss how ABM methodology can contribute, through a learning-by-modeling process, to increase our understanding of complex phenomena. We will make use of a model of university–industry relationships (UIRs) in the biotech and pharmaceutical sectors that focuses on interactions and the underlying knowledge dynamics between heterogeneous agents involved in the research process. A detailed description of this model and of the simulation results can be found in Triulzi et al. (2009). For the purpose of this article we will use this modeling experience to sustain our argumentation. Note that many other ABMs can serve the same purpose.

The article is organized as follows: We start with a conceptualization of modeling as a double-loop learning process; then a brief description of UIRs as an example of a complex phenomenon characterized by multifaceted aspects and unclear empirical evidence will be presented. This prepares the ground for claiming the necessity of a more *complex-friendly* way of analysis: ABM. Then we provide a concise description of the agent-based UIRs model and introduce the main findings. Finally, a collection of insights from the modeling experience will be summarized. We conclude by summarizing how agent-based UIRs model can be viewed as an example of double-loop learning-by-modeling process.

MODELING AS A DOUBLE-LOOP LEARNING PROCESS

ABM can be thought of as a laboratory for theory improvement. Federici et al. (2006) correctly pointed out that one of the main advantages of ABM, compared to analytical modeling, is that "computer simulations are also considered as virtual experimental laboratories to study phenomena that are difficult to observe directly" (p. 144). As Deichsel and Pyka (2009) and Brenner and Werker (2007) argued, abductive methodologies, which consist of starting by studying the facts, devising simple assumptions, and then going back and forth between assumptions and their implications, are perfectly suited for a refinement of theory. ABM contributes to improving the understanding of complex phenomena because the modeling experience is itself a learning process. More precisely, modeling is intrinsically a double-loop learning process. Indeed, as argued by Argyris and Schön (1996), two types of learning can be distinguished: single- and double-loop learning. The difference concerns the refinement of the theories and the assumptions upon which learning is grounded. The former is based on a mechanism of control and correction. When the outcome of an action differs from its expected result, a process of search and correction of the causes of the mismatch begins. This process is based on feedbacks that are analyzed using the basic knowledge of the actor, which, in single-loop learning, is an instrument of the analysis but is not itself the object of learning. Contrariwise, double-loop learning does not consider basic knowledge as a dictum. Hence, it includes all kinds of involved knowledge. Double-loop learning is suited for the investigation of complex phenomena because basic knowledge improvements are foreseen. Without doubt, this is the case for modeling experiences. Theoretical elements must be included in ABMs as well as a general awareness of the stylized facts related to the research object. In the modeling process, all of these elements are subject to validations that ultimately contribute to advance our knowledge on the issue.

AN EXAMPLE: AN AGENT-BASED MODEL OF UNIVERSITY–INDUSTRY RELATIONSHIPS

To sustain our reasoning, we choose as an example for modeling experience our model of UIRs in the biotech and pharmaceuticals sectors. A brief description of the model and its context help to familarize the reader and provides a better understanding of the peculiarities of the double-loop learning-by-modeling process.

A Complex Phenomenon: University–Industry Relationships

The way innovation is pursued in life sciences industries is deeply influenced by the advent of the biotechnology paradigm and the spread of university patenting. In this industry, UIRs have emerged as a major platform for knowledge exchange and innovation. Despite their dramatic expansion and the growing awareness in the economic and medical literature, there is only inconsistent and incomplete evidence with respect to the long-run effects on the innovativeness of the research system.

On the one hand, many authors (e.g., Blumenthal et al. 1996; Geuna 2001; Angell 2004) have suggested that UIRs can potentially damage the long-run innovativeness of the research system in life sciences. Following this literature, UIRs have modified the reward system for academic researchers, introducing a personal and institutional incentive to conduct more applied research. The possibilities to increase industry funding stemming from the commercialization of academic research potentially push universities away from pure basic research in favor of more applied research, in order to increase the probability of winding up with patentable research outcomes. The consequences of this action might be harmful for the system because it generates a situation in which fewer scientists and academic institutions are engaged in basic research, which is a fundamental component of the whole system and necessary to continuously generate innovations.

On the other hand, many studies (e.g., Meyer-Krahmer and Schmoch 1998; D'Este and Patel 2007) have shown that an incentive for university researchers to interact with industry is access to additional financial resources and industry skills. Some authors (Markiewicz and DiMinin 2004; Breschi et al. 2007; Azoulay et al. 2009) postulated the existence of a resource effect: interactions with industry, providing larger cognitive and financial resources, increase the productivity of academic scholars and university institutions in terms of publishing and patenting, thus increasing their visibility, fame, and reputation. A self-reinforcing virtuous circle is generated that ultimately boosts research productivity.

This contradictory evidence worsens because there are few empirical studies that take into consideration the system that surrounds these interactions. The role of other actors, like the government, is largely neglected. So is the bidirectionality of the knowledge and technology transfer between universities and industry. Even if there are several studies focusing on the nature of public research funding, this issue has been considered in isolation. Studies have produced inconsistent evidence: some highlight complementarities between public and private R&D, whereas others claim a substitutive relationship (for a comprehensive literature survey, see David et al. 2000).

UIRs are clearly a complex phenomenon. Therefore, these relationships call for a more comprehensive analysis that attempts to integrate the various dimensions and extend our current knowledge.

In Praise of ABMs

The lack of a generally accepted evaluation of UIRs can be traced back to two shortcomings: (1) the complexity of these relationships makes it difficult to analyze their multiple correlated effects. However, the efforts to consider the role of all the actors that are engaged in the biotech and pharmaceuticals' innovation systems and to analyze how these relationships affect each other (i.e., universities, industry, and governments) are worth making. To focus only on a selected group (e.g., universities or firms) is misleading. (2) Traditional scientific tools have only limited possibilities to disentangle the underlying complexity. Qualitative (interviews, case studies, theoretical studies, etc.) as well as quantitative analysis (surveys, econometric models, time series, etc.) are extensively applied to study UIRs. They are extremely useful to understand specific aspects of UIRs and to gain useful insights into the complex relations. But they all miss a key issue; namely, the interaction sphere dealing with knowledge flows between different actors.

Hence, we are in one of the frequent scenarios in which the complex nature of the phenomenon under investigation leads to unclear empirical evidence. How to shed light on this ambiguity?

We claim that UIRs are driven by the need to access and exchange specialized and generic knowledge by different actors. In a science-based and knowledge-intensive sector, like biopharmaceuticals, actors are heterogeneously specialized in a relatively narrow knowledge space. This is true for firms as well as for universities (though slightly less strict). The cumulative nature of knowledge in these fields leads to the creation of a self-reinforcing mechanism between, for instance, the accumulation of knowledge and expertise and the generation of successful innovations. These mechanisms are generally treated as a nuisance causing simultaneity or heteroscedasticity problems. Obviously, when it comes to the empirical analysis of evolutionary processes based on knowledge dynamics, a major problem as expressed by Keith Smith emerges: "Neither learning nor the capabilities which result, seem to be measurable in any direct way" (2005, p. 151). As shown by several models belonging to the SKIN (Simulating Knowledge dynamics in Innovation Networks) family (among others, Gilbert et al. 2001, 2007), knowledge dynamics can be effectively analyzed through multi-agent simulations based on interactions between heterogeneous and bounded rational agents. These models show that knowledge dynamics have to be placed centrally because they are the origin of agents' successes/failures on a micro level and key to understanding causes and consequences of aggregate phenomenon at the macro level. We argue that ABM's main strength is to allow a more comprehensive view of knowledge dynamics that allows for important additional insights. Obviously, traditional analytical tools are not discarded. On the contrary, they provide stylized facts and contribute to theory formation. ABMs build on them and allow going one step further: they integrate traditional analyses and provide the prerequisites for substantial theory improvement.

The Model

The model is a multi-agent simulation that reproduces R&D and knowledge dynamics in the biopharmaceutical sector, with a particular focus on the role of UIRs. We refer to a model of innovation networks originally developed by Gilbert et al. (2001, 2007). This model is further refined in subsequent works in which it has been applied to study a variety of issues related to knowledge dynamics, learning, and collaboration between agents. We extend the original model to reproduce the research environment of biopharmaceutical industries, explicitly taking into account different classes of agents moved by diverse aims and rewards (universities, biotech and pharmaceutical firms), multiple channels of interactions (research collaborations, licensing and sponsored research), and different research outputs (three classes of patents and drugs). The goal of the model is to analyze knowledge dynamics between the actors and to test the effects of agents' interactions on their knowledge base and, ultimately, on the innovativeness of the research system with the help of simulation experiments.

Types of agents

The model's population is composed of universities (UNIs), large diversified firms (LDFs), and dedicated biotech firms (DBFs). There are two further actors, a national research agency (NRA) and venture capitalists (VCs). These latter agents are funding actors (of universities and biotech firms, respectively) that are not actively engaged in research. With their research efforts, agents follow different aims. However, all firm and university actors undertake research and want to produce the best research outcomes.

Agents differ according to their knowledge base. The model's representation of the knowledge base of agents draws on the concept of *kene* developed by Gilbert (1997) and applied in previous simulations of knowledge dynamics in innovation networks. The knowledge base of each agent, its kene, consists of a vector containing different units of knowledge called *quadruples*. Each quadruple includes a research direction (RD), which allows differentiation between universities (mainly engaged in basic research) and firms (mainly engaged in applied research); a capability (C), which stands for the particular technological discipline in which actors are engaged (pharmaceutical or biotechnology); an ability (A), which reveals the actor's specialization in his or her field of capability; and an expertise (E), which shows how long an agent has been active in a certain ability.

At the beginning of each simulation experiment, the model proceeds by setting up the agents' kenes. For every agent class, specific rules have been defined. These rules allow distinguishing between different agents' classes while maintaining a fair level of intraclass heterogeneity. This is indeed one of the most significant strengths of the ABM approach: it is possible to define different classes of actors without losing heterogeneity even within each class; in other words, agent-based models allow handling a sort of squared heterogeneity. In our model we define different thresholds regarding the distribution of starting research directions (RD) and capabilities (C) among the different actors, whereas the initial abilities and expertise levels are less strictly (i.e., more randomly) defined. The number of kene elements that each actor has is proportional to his size (the size of his capital stock). However, some differences among the classes of actors exist. For instance, UNIs have a larger number of kene elements—that is, a larger knowledge base—than companies, *ceteris paribus*. A detailed table explaining kene initialization for each class of agents can be found in the Appendix.

Interactions and decision mechanisms

Agents have to make two important decisions. First, they have to allocate their funds between the different research-related activities. The set of activities that can be undertaken differ between the agents in order to reflect various interests and aims. Universities and DBFs can undertake their own research projects and joint research projects and can be active agents (licensors) in licensing agreements. University research can be sponsored by an LDF, whereas this option is not available to DBFs. In this case, funding for a university research project comes from an LDF, which, in turn, acquires intellectual property rights related to the project's results under the form of a royalty-free license. LDFs can undertake own research or joint research with DBFs, buy licenses from universities or DBFs, and sponsor university research. Budget allocation over the set of research activities follows a satisfycing behavior similar to Nelson and Winter's (1982) evolutionary model. Accordingly, agents initially decide how to allocate their budget randomly. When firms are successful, they stick to the same allocation strategy. When firms are not successful, they modify their strategy.

If agents decide to allocate part of their resources to joint research projects, the follow-up step is the choice of one or more partners with whom they cooperate in a project. This is made according to two different partnership strategies: a conservative strategy, which aims to find—concerning their knowledge bases—similar agents, and a progressive strategy, which aims to find different partners (according to an agent's absorptive capacity; see also Pyka et al. 2007). The actors that choose the former strategy aim to undertake an incremental research project; therefore, they prefer less risk and a common understanding. In this case, the variance between the knowledge bases of the partners is small; this increases the probability of success but, on the other hand, reduces the potential magnitude of the project outcome. Instead, actors that choose to follow a progressive strategy aim to undertake a radical research project. In this case, the variance between the knowledge bases is high, with a positive effect on the potential innovativeness of the project outcome but with a negative effect on the probability of success. Agents also weigh their partners; this means that they will first look to previous partners with whom they have conducted a successful research project in the past.

The environment

The model's environment plays an important role: agents are aware of competition as well as the possibility to cooperate. Both options affect firms' decisions.

Firms observe and evaluate competitors' behaviors when they decide on their allocation strategy between different research-related activities. Periodically firms compare their allocation strategy with the average allocation per activity of the most successful firms; that is, those in the first quartile of firms ranked according to their capital stock belonging to the same class (DBFs or LDFs). If a firm is successful (meaning that it is in the first quartile), it does not change its allocation strategy. Firms that are not successful gradually change their strategy, imitating successful firms.

Agents can also opt for cooperation with other agents. The degree of inter-agent cooperation again is influenced by the environment. A successful firm that is extensively engaged in cooperation acts as a signal to other firms that will urge them to adopt the same strategy. What makes collaboration a factor of success is the knowledge exchange that comes with it. Indeed, as in Gilbert et al. (2001), the process of one's own and collaborative research is based on the combination of selected elements of an agent's knowledge base, which forms a so-called innovation hypothesis (IH). In the case of joint research, the project knowledge base is a combination of parts of the knowledge bases of the involved agents. Some quadruples of the agents' kenes are randomly recombined to form a project innovation hypothesis. If the project is successful, actors with an absorptive capacity (Cohen and Levinthal 1990) above a critical threshold acquire the knowledge of the joint innovation hypothesis that has been contributed by the project partner(s), though with a reduced experience level. This enables agents that are engaged in collaboration to upgrade and expand their knowledge base.

Model's dynamics

Each simulation run consists of several iterations; that is, cycles of research. A cycle starts when actors choose to start a joint research project, their own research project, or both. In the former case, the actors look for partners and subsequently jointly run the project (and share the project costs and the ownership of the outcomes). In the latter case, the actors set up and run the project in isolation. The project lasts several periods and finally is evaluated. If the project is successful, a patent is granted. There are three kinds of possible outcomes in the model, both for one's own and joint research projects (ranked from the most to the least innovative): (1) A-class patent, (2) B-class patent, and (3) C-class patent. Which outcome is generated depends on the research direction of the project (a basic research direction increases the likelihood of obtaining an A-class patent) and on the variance of the involved capabilities (the higher the variance, the higher the outcome's value). The probability of success, which is positively related to the agent's experience level, is higher for an applied research direction and negatively depends on the variance of the capabilities involved.

If the patent is granted to a university or a DBF, the patent holder enters the market for research and tries to find an LDF willing to acquire the license and to pay the related royalties. If the patent is originally granted to an LDF, the firm can directly conduct clinical trials and try to develop a new drug to earn revenues. Eventually, the actors reinvest the money that they have gained at the end of the research cycle (from the license's royalties or the sales of the new drug) in new research projects and a new cycle begins.

RESULTS

Several Monte Carlo simulation experiments based on a standard and some alternative scenarios were performed. Results of different scenarios are compared and the appropriate tests of statistical significance are performed.

Results are shown in Figure 1. Panel A shows the dynamics of the average research orientation of universities for two simulation experiments: a standard scenario in which universities are allowed to interact with industry (both DBFs and LDFs) and a second one in which universities were the only agents in the simulation population. One immediately notices that in the latter case universities maintain a strong focus on basic research (lower values of the average research direction). This shows that relationships with industry do increase incentives for universities to engage in applied research.

Panel B shows the results for the test of the so-called resource effect hypothesis. On the vertical axis we find the percentage of innovative patents (A-class) generated by universities relative to total university patenting. Again, two simulation experiments were performed. In the case of interactions between universities and DBFs (standard scenario), university patenting was more innovative (larger percentage of A-class patents) than in the case in which these interactions were not permitted (without_DBFs scenario). This finding shows that universities do not enjoy cognitive resource effects related to interactions with industry. We also tested the hypothesis of a financial resource effect generated by licensing revenues from LDFs. The results in terms of the relative number of A-class patents did not prove to be statistically significant. However, we found that interactions with LDFs do increase the total number of university patents but without influencing their innovation value.

We also tested the effects of interactions with universities on DBFs' innovative capabilities. The percentage of innovative patents (A-class) out of the total number of DBFs' patents can be found on the vertical axis of Panel C in Figure 1.

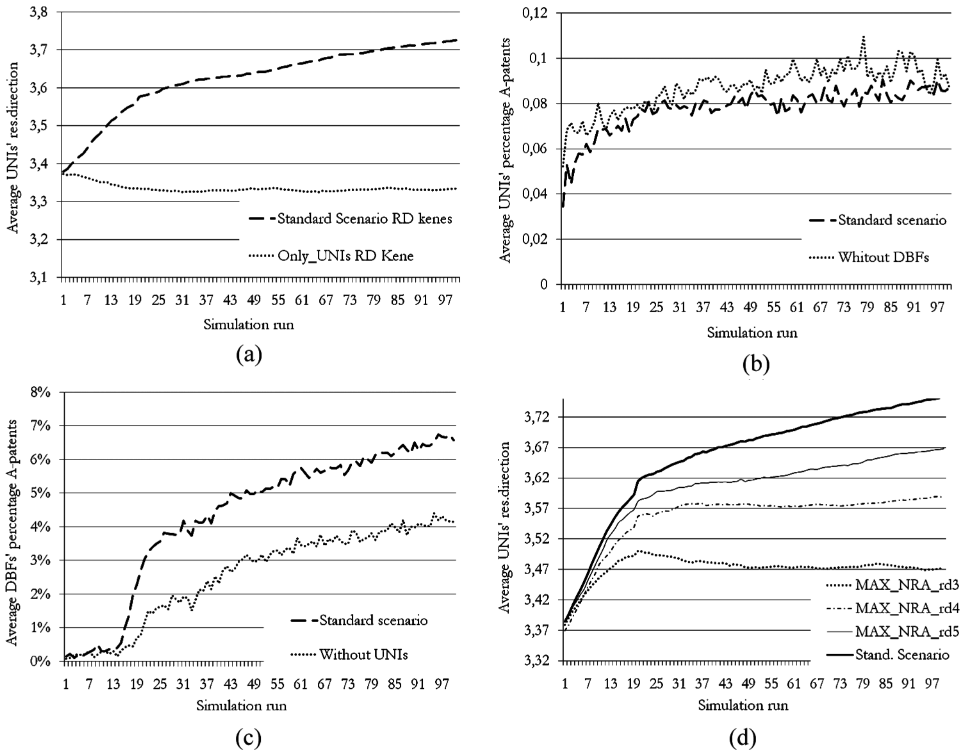

FIGURE 1 Comparison between the different scenarios' results.

The two lines represent the values of this percentage in two different simulations: the standard scenario and a scenario in which universities were excluded from the population (DBFs could not interact with them). The difference in the trends shows that DBFs greatly benefit from knowledge exchange with universities. This shows that despite the missing cognitive resource effects on the universities' side, this effect is visible for the industry partners. In other words, biotech firms benefit more than universities from the knowledge exchange.

Finally, our simulation experiments showed that governments can effectively reduce the harmful effects of UIRs on the universities' research direction. Higher incentives for basic research can be restored through a larger basic research public funding budget. This is shown in Panel D in Figure 1. On the vertical axis we find the average research direction of universities. The different lines represent different scenarios in which the basic research public budget is progressively increased (MAX_NRA_rd5/4/3). The graph shows that a larger government basic research funding prevents the shift of the research orientation of the universities from basic to applied research.

INSIGHTS FROM THE MODELING EXPERIENCE

The application of agent-based simulation methodology generates new insights into the complex phenomenon of university–industry relationships.

Our experiments reject the hypothesis of a positive influence of a cognitive resource effect on university innovative patent productivity and partially reduced the influence of the financial resource effect. Nevertheless, knowledge exchange processes are still a crucial characteristic of UIRs. Intensive interactions between universities and DBFs do not seem to improve universities' innovative capabilities. This, however, does not mean that there are no benefits from these interactions. Through joint research projects, an exchange of knowledge between universities and biotech firms occurs. These newly acquired capabilities and abilities expand the university knowledge base, thereby increasing its heterogeneity, but this is not sufficient to improve universities' patent productivity. Many universities might not have the right skills to deal with applied research; that is, they might not be experienced enough to deal with a knowledge that is far from their traditional research orientation. This effect is due to the fact that universities acquire the kene elements that the DBF partner has contributed to the knowledge base of the joint research project, but these kene elements add to a university's knowledge base with a basic level of experience only. This negatively affects the probability that a following project using the newly acquired knowledge generates an innovative outcome. A certain time lag is required to master the new knowledge; hence, several attempts are required before the new knowledge starts to be productively used in the following projects.

Our findings also show that when a complex phenomenon like UIRs involving heterogeneous actors is analyzed, one has to consider all of its multifaceted aspects. In particular, our results show that UIRs cause a significant increase in the innovative potential of biotech firms. This is due to a threefold effect. First, interactions with universities expand DBFs' knowledge bases, allowing biotech firms to absorb new kene elements focusing on fundamental research. Second, they also increase the variance in their capabilities and, third, they have a positive effect on DBFs' networking experience. Therefore, our results highlight the importance of UIRs concerning technology and knowledge flows. These finding suggest that, in addition to universities, new scientific knowledge is increasingly generated by biotech firms.

Finally, our results show that even if public and industry research funding are sometimes seen as substitute, in the biotech and pharmaceuticals they are complementary. According to our results, government basic research grants are crucially important to counterbalance the different aims and incentives provided by industry, which further enlarge the market failure, especially in the long run. Accordingly, government research policies should be oriented to raise the public research funding budget with the aim to ensure that an adequate amount of fundamental research is undertaken by universities.

DISCUSSION AND CONCLUSIONS

Today innovation activities are undertaken in extremely complex systems, which are characterized by heterogeneous actors, multidimensional interactions, and multiple knowledge flows. The increasing complexity creates new challenges for scholars and can be disentangled only by the integration of several methodologies.

We argue that ABM can play a key role in this respect. In particular, the double-loop learning-by-modeling process challenges what we think we know about the subject of the analysis. An example of this double-loop learning process focused on UIRs has been provided. As argued by Argyris and Schön (1996), double-loop learning starts from theory and refines it through a verification and validation cycle in which the theory itself is checked and improved. In a similar vein, Deichsel and Pyka (2009) recommended that modelers should start with reasonably simple assumptions, based on theoretical elements and stylized facts, and hence enter the process of going back and forth between assumptions and implications. When the researcher faces a complex phenomenon with multi-agent and multidirectional interactions, a broad view is necessary. Nevertheless, understanding what is driving the observed phenomena requires identification of the different components of the system and how the various channels of interactions affect them. Putting these components together and using them in a modular way (see Mallavarapu et al. [2009] for an example in systems biology) allows grasping how the system works as a whole. Such a modeling exercise shows us what we do not know about the inter-component interactions and automatically guides the direction of future research.

Figure 2 graphically illustrates our definition of a double-loop learning-by-modeling process, as adapted from Argyris and Schön's (1996) original formulation.

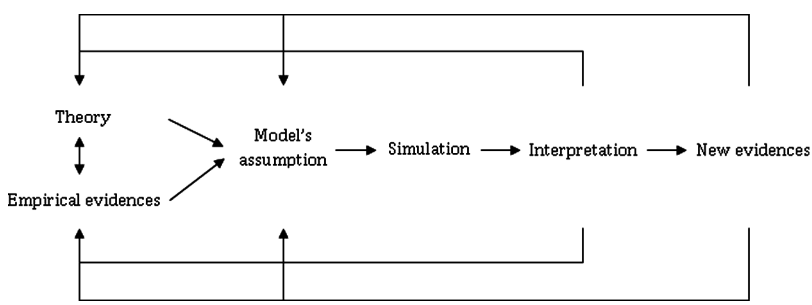

FIGURE 2 Double-loop learning-by-modeling process.

The reader who is familiar with the literature discussing the relation between multi-agent-based simulations (MABS) and multi-agent systems (MAS) may have noticed that the concept of double-loop learning by modeling also builds on the works by Edmonds (2001) and Federici et al. (2006). The former defines the basic modeling relation as one in which the model goes through a process of encoding a natural process (i.e., abstracting the dynamics of the targeted system into an MAS), inference based on MAS simulations, and decoding (analysis of results and interpretation back onto the targeted system). Edmonds also pointed out that: "the above characterisation of the MABS modeling process is simplified, in part because only rarely is direct modeling attempted" (p. 271). The author argued that there are often several other layers of models between the targeted system and the MAS. Federici et al. highlighted this aspect and specified that "in a MABS reality is not directly mapped in a MAS [. . .] but what that is mapped it is a model of reality. This model is an intermediate step between reality and MAS" (p. 146). The double-loop learning by modeling methodology that we have discussed in this work provides a helpful framework to deal with the interactions between reality and modeling. Evidently it is not possible to fully understand all of the interactions, causal relations, and variables that affect a complex system in reality (if we could we would not have any research question). Hence, it is necessary to start from an approximation of reality that is based on the current state-of-the-art knowledge of the targeted system provided by empirical evidences, stylized facts, and theories. Assumptions upon which the model is defined draw on these building blocks. The state-of-the-art knowledge is further improved because the interpretation of the simulations results during both the intermediate steps of testing the model and the analysis of its final results can generate new evidence that, through a double-loop learning process, can lead to new theoretical and empirical insights. Of course, the contribution of the model is greater the larger the gap between reality and the state-of-the-art knowledge. Theories, empirical evidence, and MABS have to build on each other. Paraphrasing Edmonds (2001), we argue that theories and evidence help understanding what the truly relevant aspects of the targeted system's structure to be included in the MAS are. Furthermore, it is important to consider theories and evidence as flexible scientific constructs that can be changed and improved in the interpretation step of the modeling process.

In the case of the UIRs model, we used as initial inputs of the modeling process a combination of theories and insights gained from some empirical evidence as well as some stylized facts about UIRs. In particular, we started from the evolutionary process of innovation and technological change as originally expressed by Nelson and Winter (1982). These theories greatly help to set up a more realistic view of agents (bounded) rationality (Simon 1958) and of their satisfycing behavior (for an extensive overview on theoretical approaches to tackle learning in an evolutionary environment, see Dosi et al. [2005]). Our modeling effort also benefited from the large literature about the nature of innovation processes involving science and technology interactions; namely, from the dichotomy between the linear and the so-called chain-linked models of innovation (Kline and Rosenberg 1986). We also refer to theories of innovation networks (e.g., Buchmann and Pyka 2011), which are now largely accepted as a crucial source of innovation driven by interactions between specialized agents and the underling knowledge exchange (see, for instance, Gilbert et al. 2001). Finally, we also built on theories on the relation between cognitive distance and innovation outputs (Nooteboom 1992, 1999) or the relations between research orientation and innovation output to set up the equations that transform innovation hypotheses into research outcomes.

A crucial step in the double-loop learning process is the interaction between theories and evidence. We matched theories with some facts and figures about UIRs in the biopharmaceutical industry. For instance, we calibrated the functional relations between the kene-related variable (project's research direction, cognitive distance, and experience level) and the type of research outcome as well as the parameters of our model by running several test simulations and comparing their results with empirical data about the percentages of A/B/C drugs in the United States from the U.S. Food and Drug Administration. Then we picked up those parameters and we refined the equations in such a way that made the results reasonably similar to real-world data. Our model in its present stage does not claim to reproduce the developments in the biopharmaceutical industries in a history-friendly way (Malerba et al. 1999). Instead, we focus on important stylized facts and elaborate the complex interaction patterns among the various actors. In this sense, our modeling approach claims to be empirically guided. For an extensive overview on the relationship between ABM and empirical evidence, see Fagiolo et al. (2007). Finally, the robustness of our parameters has been tested through traditional sensitivity analysis. Surprising findings, compared to the initial expectation of the modeler, can emerge. This eventually leads to an interpretation of the results that really provides new insights and viewpoints. In our model, this is the case, for instance, for the lack of experience of many universities in dealing with a too radically different research orientation compared to their traditional one. Moreover, additional research questions came out of the double-loop learning-by-modeling mechanism. During the modeling process, we realized that there are some important elements that can affect the long-run innovativeness of the research system in the UIRs framework that were not initially considered and are not fully explained in the literature. These elements include (1) how changes in the budget allocation policies of the NRA (for instance, toward a stronger selectivity of the funding recipients) influence universities' behaviors with respect to interactions with industry and how this affects the patterns of interaction between the agents; and (2) how changes in the agents' investment strategies due to, for instance, external shocks affect the innovativeness of the research system through the emergence of different UIRs patterns.

A common objection to this methodology is that results still depend on assumptions. In our case what we say is that *if we assume* bounded rational agents and such particular relations between research orientation and innovative outputs and between cognitive distance and innovative outputs, then university–industry relationships lead to the knowledge dynamics that we have highlighted. We agree with Edmonds and Bryson (2004) when they urged to "seek scientific foundations for agent systems." Indeed, concerning our work, it is crucial to notice that the above-mentioned assumptions have not "fallen from the sky" but are a result of a sort of inductive process driven by the interaction among theoretical elements, stylized facts, and learning-by-modeling. The value-added of the double-loop learning mechanism in our particular case is that, on the one hand, analyzing the effects of UIRs on the long-run innovativeness of the research system considering only bilateral interactions between the agents would have neglected the cumulative effects of multidirectional inter agent interactions. On the other hand, the theories explaining these effects have also been critically revised during the modeling process, and alternative explanations of the interactions between the different elements emerged. Thus, the ABM methodology can substantially contribute to a better understanding of complex socioeconomic interactions and thus support the development of theories that are suited to dealing with this complexity without ignoring it.

REFERENCES

Angell, M. *The Truth About Pharmaceutical Companies: How They Deceive Us and What to do About it*. New York: Random House, 2004.

Argyris, C. and Schön, D. A. *Organizational Learning: A Theory of Action Perspective*. Reading, MA: Addison-Wesley, 1996.

Azoulay, P., Ding, W., and Stuart, T. "The Impact of Academic Patenting on the Rate, Quality and Direction of (Public) Research Output." *The Journal of Industrial Economics* 57, no. 4 (2009): 637–76.

Blumenthal, D., Causino, N., Campbell, E., and Louis, K. S. "Participation of Life-Science Faculty in Research Relationships with Industry." *The New England Journal of Medicine* 335, no. 23 (1996): 1734–39.

Brenner, T. and Werker, C. "A Taxonomy of Inference in Simulation Models." *Computational Economics* 30 (2007): 227–44.

Breschi, S., Lissoni, S., and Montobbio, F. "The Scientific Productivity of Academic Inventors: New Evidence from Italian Data." *Economics of Innovation and New Technology* 16, no. 2 (2007): 101–18.

Buchmann, T. and Pyka, A. "Innovation Networks." In *Handbook on the Economics and the Theory of the Firm*, edited by J. Krafft and M. Diettrich. Cheltenham, UK: Edward Elgar, forthcoming.

Cohen, W. M. and Levinthal, D. A. “Absorptive Capacity: A New Perspective on Learning and Innovation.” *Administrative Science Quarterly* 35, no. 1 (1990): 128–52.

David, P. A., Hall, B. H., and Toole, A. A. “Is Public R&D a Complement or Substitute for Private R&D? A Review of the Econometric Evidence.” *Research Policy* 29 (2000): 497–529.

Deichsel, S. and Pyka, A. “A Pragmatic Reading of Friedman’s Methodological Essay and What it Tells Us for the Discussion of ABMs.” *Journal of Artificial Societies and Social Simulation* 12, no. 4 (2009).

D’Este, P. and Patel, P. “University–industry Linkages in the UK: What are the Factors Underlying the Variety of Interactions with Industry?” *Research Policy* 36, no. 9 (2007): 1295–1313.

Dosi, G., Marengo, L., and Fagiolo, G. “Learning in Evolutionary Environments.” In *Evolutionary Principles of Economics*, edited by K. Dopfer, 255–338. Cambridge, UK: Cambridge University Press, 2005.

Edmonds, B. “The Use of Models: Making MABS More Informative.” In *Lecture Notes in Computer Science*, edited by S. Moss and P. Davidson, Vol. 1979, 269–82. Boston, MA: Springer, 2001.

Edmonds, B. and Bryson, J. “The Insufficiency of Formal Design Methods—The Necessity of an Experimental Approach for the Understanding and Control of Complex MAS.” In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, Vol. 2, 938–45. New York: IEEE Computer Society, 2004

Fagiolo, G., Birchenhall, C., and Windrum, P., eds. “Special issue on “Empirical validation in agent-based models.” *Computational Economics* 30, no. 3 (2007).

Federici, M. L., Redaelli, S., and Vizzari, G. “Models, Abstractions and Phases in Multi-agent Based Simulation.” In *Proceedings of the 7th WOA Workshop: From Objects to Agents*, edited by Andrea Omicini and Mirko Viroli. Catania, Italy: CEUR Workshop Proceedings, 2006.

Geuna, A. “The Changing Rationale for European University Research Funding: Are There Negative Unintended Consequences?” *Journal of Economic Issues* 35, no. 3 (2001): 607–32.

Gilbert, N. “A Simulation of the Structure of Academic Science.” *Sociological Research Online* 2, no. 2 (1997).

Gilbert, N., Ahrweiler, P., and Pyka, A. “Learning in Innovation Networks, Some Simulation Experiments.” *Physica A: Statistical Mechanics and its Applications* 374, no. 1 (2007): 100–9.

Gilbert, N., Pyka, A., and Ahrweiler, P. “Innovation Networks–A Simulation Approach.” *Journal of Artificial Societies and Social Simulation* 4, no. 3 (2001).

Kline, S. and Rosenberg, N. “An Overview of Innovation.” In *The Positive Sum Strategy: Harnessing Technology for Economic Growth*, edited by R. Landau and N. Rosenberg, 275–305. Washington, DC: National Academy Press, 1986.

Malerba, F. “Learning by Firms and Incremental Technical Change.” *The Economic Journal* 102 (1992): 845–59.

Malerba, F., Nelson, R., Orsenigo, L., and Winter, S. “History-friendly Models of Industry Evolution: The Computer Industry.” *Industrial and Corporate Change* 8 (1999): 3–40.

Mallavarapu, A., Thomson, M., Ullian, B., and Gunawardena, J. "Programming with Models: Modularity and Abstraction Provide Powerful Capabilities for Systems Biology." *Journal of the Royal Society Interface* 6 (2009): 257–70.

Markiewicz, K. and DiMinin, A. "Commercializing the Laboratory: The Relationship between Faculty Patenting and Publishing." Working paper series IN-SAT Lab, Scuola Superiore Sant'Anna, 2005. (WP-2004/02)

Meyer-Krahmer, F. and Schmock, U. "Science-based Technologies: University–Industry Interaction in Four Fields." *Research Policy* 27, no. 8 (1998): 835–51.

Nelson, R. R. and Winter, S. *An Evolutionary Theory of Economic Change*. Cambridge, MA: Harvard University Press, 1982.

Nooteboom, B. "Towards a Dynamic Theory of Transactions." *Journal of Evolutionary Economics* 2 (1992): 281–99.

Nooteboom, B. *Inter-Firm Alliances: Analysis and Design*. London: Routledge, 1999.

Pyka, A. and Fagiolo, G. "Agent-based Modelling: A Methodology for Neo-Schumpeterian Economics." In *Elgar Companion to Neo-Schumpeterian Economics*, edited by H. Hanusch and A. Pyka, 467–87. Cheltenham, UK: Edward Elgar, 2007.

Pyka, A., Gilbert, N., and Ahrweiler, P. "Simulating Knowledge Generation and Diffusion Processes in Innovation Collaborations and Networks." *Cybernetics and Systems* 38, no. 7 (2007): 667–94.

Simon, H. A. "Rational Choice and the Structure of the Environment." *Psychological Review* 63 (1956): 129–38.

Smith, K. "Measuring Innovation." In *Oxford Handbook of Innovation*, edited by J. Fagerberg, D. C. Mowery and R. R. Nelson. New York: Oxford University Press, 2005.

Tesfatsion, L. and Judd, K. L., eds. *Handbook of Computational Economics: Agent-based Computational Economics*. North Holland, The Netherlands: Elsevier, 2006.

Triulzi, G., Scholz, R., and Pyka, A. "R&D and Knowledge Dynamics in University–Industry Relationships in Biotech and Pharmaceuticals: An Agent-based Model." FZID discussion paper series, University of Hohenheim, Stuttgart, Germany (forthcoming, 2011).

World Bank. *World Development Report*. Washington, DC: Oxford University Press, 1999.

APPENDIX

TABLE A1 Initial Distribution of Agent's Kene Elements

Agents	Kene quadruple's element	Distribution
UNIs	RD	70% of UNIs'kene quadruples have a research direction randomly distributed within the range 1–3
		30% of UNIs'kene quadruples have a research direction randomly distributed within the range 2–7
	C	80% of UNIs'kene quadruples are distributed around 3–5 capabilities' focal points;the rest arerandomly spread
DBFs	RD	10% of DBFs'kene quadruples have a research direction randomly distributed within the range 1–4
		80% of DBFs'kene quadruples have a research direction randomly distributed within the range 3–6
		10% of DBFs'kene quadruples have a research direction randomly distributed within the range 5–9
	C	90% of DBFs'kene quadruples are randomly distributed within the range of capabilities from 1 to 60; 80% of this 90% is distributed around a $\pm$ 5% range of the first capability
		10% of DBFs'kene quadruples are randomly distributed over the whole range of capabilities (1–100)
LDFs	RD	20% of LDFs'kene quadruples have a research direction randomly distributed within the range 1–9
		10% of LDFs'kene quadruples have a research direction randomly distributed within the range 2–4
		70% of LDFs'kene quadruples have a research direction randomly distributed within the range 5–9
	C	60% of LDFs'kene quadruples are distributed within the range of capabilities from 61 to 100 with focal points with group 3–5 kene elements
		40% of LDFs' kene quadruples are distributed over the whole range of capabilities (1–100) with focal points with group 3–5 kene elements