



LABORatorio R. Revelli
Centre for Employment Studies

The Methodological Status of Agent-Based Simulations

Roberto Leombruni
LABORatorio Revelli
SIEC – University of Ancona

Working Papers Series
No. 19

Abstract

In this paper I give a tentative assessment of the methodological status of agent based simulations. I first show under which conditions ABS can be a complement to traditional modelling. I then consider whether they can be held as a sound methodology of their own. Various topics relevant to the argument are briefly discussed, such as the forecasting role of theories, the realism of assumptions, Hayek's insights on economics methodology. I cast the arguments given into some results of modern theory of abductive inference, providing a framework that helps clarify the logical status of simulations, and gives some hints on how to foster their role as a self sustained tool for economic reasoning.

1 Introduction

According to Hempel, scientific explanation of a phenomenon consists in its reconduction to a general law¹. Monetarist explanation of inflation is an example, in that it ascribes inflation to the "law", stating that every monetary expansion will end up in a rise of the aggregate price index. In this view, the core of scientific endeavour is the discovering of new laws, the building of new theorems.

In this view, agent-based simulations can not stay at the core of scientific endeavour, since they do not provide us with theorems. The generalisation of their results, indeed, conflicts with the impossibility of spanning numerically the domain of all relevant variables and parameters. In a sense, a simulation provides us with just anecdotic evidence².

Reducing our ambitions, we could find a role for simulations in their use as a laboratory. Social scientists sometimes complain the fact that they can not test hypotheses under controlled conditions. Could simulations do that job? Strictly speaking, one should answer *no*. To bring a consumer into a laboratory to study its behaviour is a matter of experimental economics: Into a personal computer you can bring just the *model* of a consumer.

Let us reduce again our ambitions, and ask ourselves whether a simulation can be seen as a mean to explore in a laboratory-like manner some field of interest, using models as a surrogate of the actors involved. This time the answer is *yes*. It should be noted, however, that this is the same kind of activity a researcher does when he takes his pencil and tries to put together various equations — various models of the actors he is studying. With two differences: handling a pencil is easier, and if you have the good luck of arriving to a nice result, that result will resemble to a theorem, to a "piece" of science.

Certainly, if this were the "true" picture, it wouldn't be very encouraging. Fortunately, the argument sketched is not that sound,

¹The so called deductive-nomological, or "covering-law" view of scientific explanation. The main references are Hempel and Oppenheim (1948) and Hempel (1965, 1966). Deductive-nomological explanations of an explanandum E are "deductive arguments whose conclusion is the explanandum sentence E , and whose premiss-set, the explanans, consists of general laws and of other statements which make assertions about particular facts" (Hempel 1966, p. 51).

²For the sake of readability I will hereafter use "simulation", "agent-based simulation" and "agent-based model" interchangeably.

mainly in one point: It rests of a view of scientific explanation that — particularly in social sciences — calls for many refinements³.

As a matter of fact, agent-based simulations are diffusing in many disciplines, from biology, to chemistry, to economics. Provided that this, in turn, is not a sound evidence of their value, I will give in this paper some arguments on the importance they can have in economic reasoning.

To set up the stage, I will start with few words on the origin and definition of agent-based simulations. This will help me to assess a first case — a "light" one — in which simulations provide a sure added value to the scholar.

I then proceed showing why agent-based computational economics (ACE) has to be considered as a sound methodology of its own. The line of reasoning has three steps. First, I will clarify (one of) the arguments upon which the covering-law view of scientific explanation misses some specificities of economic research. Next, I will show how ACE can be useful in tackling with these specificities. Finally, I will generalise the intuitions given in first two steps, proposing a tentative assessment of their methodological status.

2 Basic features of agent-based modelling and computation

Agent-based computational economics is a recently established field of research whose lineage is twofold. From one side, it is a derivation of computational economics, quite similar in its philosophy to the pioneering studies on microsimulations. Epstein and Axtell, tracing back the origins of the agent-based approach in social sciences, quote an interesting passage of Orcutt *et al.*, that at the beginning of the Sixties defined this way their main concern:

Our socioeconomic system is a complicated structure containing millions of interacting units, such as individuals, households, and firms.

It is these units which actually make decisions about spending and

³As a matter of fact, the deductive-nomological view of scientific explanation today is far from being the "received view" in the philosophy of science literature. A detailed reconstruction of the debate on the "crisis" of the D-N model and on the alternative views of scientific explanation can be found in Salmon (1988). A survey of the many refinements and/or alternatives approaches that have been proposed in economics is out of my scope. Many insights, however, can be found in Antiseri (1997), where there is an up to date discussion dealing with the peculiarities of social enquiry.

saving, investing and producing, marrying and having children. It seems reasonable to expect that our predictions would be more successful if they were based on knowledge about these elemental decision-making units — how they behave, how they respond to changes in their situations, and how they interact.⁴

From the other side, ACE exploits a technology with a much more "complex" lineage, in that it combines elements of many fields ranging from cellular automata, to genetic algorithms, to artificial intelligence, to artificial life. During the last two decades, the boost in computers' power and the diffusion of programming languages with an object-oriented architecture⁵, provided the tools for a widespread exploitation of this technique within the broad field of studies usually labelled as complexity sciences.

As regards their origins, then, agent-based simulations are far from being a peculiarity of economic research — the so-called "economics imperialism" here is not in action. The same can be said about the advancements they promise. As a matter of fact, one of the most appealing features of this methodology for the social scientist is the possibility it gives of performing an interdisciplinary work (Page 1999, Axelrod 1997, Epstein and Axtell 1996).

To focus on the role of agent-based simulations in economics⁶, let us consider the following definition given by Leigh Tesfatsion (2000), one of the pioneers of ACE:

agent-based computational economics [...] is roughly defined by its practitioners as the computational study of economies modelled as evolving systems of autonomous interacting agents. A principal concern of ACE researchers is to understand the apparently spontaneous formation of global regularities in economic processes, such as the unplanned coordination of trading activities in decentralized market economies that economists associate with Adam Smith's invisible hand. The challenge is to explain how these global regularities

⁴Epstein and Axtell 1996. According to the two authors, "micro-simulation has more of a 'top-down' character since it models behavior via equations statistically estimated from aggregate data, not as resulting from simple local rules." I will get back to this point later on.

⁵Interestingly enough, the ancestor of modern object-oriented programming languages, Simula, was developed at the beginning of Sixties precisely for simulation purposes. This kind of technology revealed much powerful, and an object-oriented architecture characterises the today most diffused programming languages as Java and C++.

⁶Some notes on the specificities of simulations in economics *versus* simulations in social sciences can be found in Terna and Conte 2000.

arise from the bottom up, through the repeated local interactions of autonomous agents channeled through socioeconomic institutions, rather than from the top-down imposition of fictitious coordination mechanisms such as market clearing constraints or an assumption of single representative agents.

Forgetting for a while the "C" of ACE, the "study of economies modelled as evolving systems of autonomous interacting agents" has a long tradition in economics, being one of the main tenets — for instance — of evolutionary economics, as Tesfatsion points out in her guide to the literature. The research programme focusing on heterogeneity and interaction (H&I) has certainly many points of contact with this approach, too. The same could be said about Austrian economics. The point is that, strictly speaking, also general equilibrium theory has an "agent-based" nature. Is that "C" the real mark of the approach, then? Certainly not: The definition quoted stressed some other points that are less common to, say, general equilibrium theory. It is anyway a good point of departure.

A common practice used by scholars dealing with models whose analytical solution is not one-way, is to substitute "realistic" values in the parameters and compute them numerically. It is a way of evaluate various features of a model, to test the relevance and sensitivity of its implications, and to explore additional hypotheses⁷. With a fully-fledged ACE approach one can make this job more extensively, getting much deeper into the problem studied. The change of the functional form which defines the preferences of a consumer is as immediate as that of the value of a parameter; the substitution of that function with a procedure implementing some rule of thumb is equally immediate; heterogeneity among agents, or out of equilibrium dynamics are other topics easily handled. More in general, once an analytical model has an ACE counterpart, many of the simplifications usually made for the sake of tractability can be relaxed quite easily, to "see what happens".

The added value lies here, in the possibility of growing into a computer artificial markets where to test many extension the researcher

⁷As Judd points out (1997), even these opportunities are widely under-exploited. He refers to the traditional tools developed within computational economics, presented both as complements and as substitutes to theorem building. The arguments presented in this section partially overlap to his work: At the end, some added value stems directly from the use of a pc. In the following, however, I will focus mainly on the specificities of agent-based computation, both with respect to theorem building and to traditional computational approaches.

is interested in.

To be true, the access to these benefits is not immediate, since an ACE counterpart is not always available.

A first case to cite, has to do with the class of models where we know — by means of some fixed-point theorem — that an equilibrium exists, but where we know — by means of some complexity result — that the solution is not computable. In these cases, there is no way of implementing the agents living in such models.

A second, somewhat weaker, impossibility, has to do with optimal behaviour. The implementation of agents whose behaviour is ruled by the (computable) maximisation of an intertemporal payoff function has some drawbacks. Let us assume that the optimisation has, as a solution, a closed-form function of a given set of parameters. The simulation is feasible, but the benefits we pointed out in terms of flexibility are hampered: A change in the assumptions would typically require the re-writing of the model to solve a new maximum problem. In other words, the simulation would grant a little more than the computation of the model solution using "realistic values" for the parameters. On the other hand, when the optimisation had not a closed-form solution, the behavioural rules have to be approximated numerically. The simulation is still feasible, and more traditional tools developed within computational economics help doing this job. In an ACE framework, however, a population of heterogeneous agents performing a (asynchronous) numerical maximisation of an intertemporal payoff function can easily give rise to simulations with a cumbersome architecture, hardly providing hence easy to read results.

Someway, simulations behave *as if* they didn't like textbook models. Indeed, they are particularly suited for the study of markets populated by individuals with bounded-rationality and interacting with each other, and particularly for the study of their dynamical behaviour. As Page puts it, "Computational models enable the study of complex, dynamic worlds because themselves are dynamic" (1999).

In any case, *if* one is able to formulate a model where:

- (1) The "building blocks", individuals or institutions, are modelled as agents with a computable behaviour;

- (2) The dynamics of the model is defined from "the bottom up", thanks to a computable schedule where there is the explicit list of the actions that will "happen" in each time unit;

then, the analytical study is capable of being supported by an ACE counterpart with which to test any extension the researcher is interested in, compatible with requirements (1) and (2)⁸.

While the first requirement does not imply a particular effort, the second one is somehow at odds with more traditional modelling. It simply requires, however, to substitute (or to add) to the formal study of equilibrium conditions the writing down of the ordered list of events that will actually happen, the market being in equilibrium or not. The passage, is from the simultaneous equations style of general equilibrium theory (and of celestial mechanics) to a story telling style (quite typical in social sciences).

Summing up, ACE approach offers to the scholars, particularly — but not exclusively — to those interested in fields as bounded rationality, H&I or non linear dynamics, a powerful tool that *complements* more traditional formal reasoning with some added value in terms of model testing and extendibility. The price required to buy such benefits is that the architecture of the model has to be designed in a way compatible with its translation into a computable procedure.

3 The challenge of complexity. Hayek's explanations of the principle

Till here, I gave a "light" assessment of ACE capabilities in supporting more traditional ways of doing research.

Now, I will address the question whether ACE modelling can be considered a sound methodology of its own. We saw that a clear way to point out the possible limits of a simulative approach is

⁸A proof of this proposition would account to no more than a tautology: The definition of agents behaviour (and state variables), together with the choice of a starting point for the flow of agents' actions are the necessary and sufficient elements of an agent-based computer simulation. The computability of the simulation, i.e., the finite-time requirement for the computation of an artificial time unit, is more difficult to assess, particularly when the actions of an agent can generate a cascade of other agents' actions. Software packages like Swarm heavily simplify the project of computable simulations, since a built-in management of the schedule is provided (see www.swarm.org).

by contrasting the anecdotic nature of its outputs with the view that a "scientific" explanation of a phenomenon should rest on its reconduction to a law-like proposition. Does this latter view meet with the specificities of economic research? Indeed, the answer rests on one's opinion about the main concern of Economics.

Simplifying a bit, we can contrast two such opinions. The first one is well exemplified by Robbins' classical definition: Economics has to do with the allocation of scarce means to competing ends. The search for the optimal allocation is a built-in feature of this view; the natural toolbox is mathematics, and theorem building. The second one can be traced back to the very origins of economic research, and has had in Hayek one of its main partisans: Economics has to do with the unwanted effects of individual actions on aggregate behaviour. Germane to this view is the stress on market process, on aggregation, and on the complexity of the economic system. There is not a natural toolbox for performing this study: The historical analysis on the emergence of the QWERTY standard is a nice example of a not "theorem-oriented" one⁹.

As regards methodology, the first view has a benchmark in Friedman's *The Methodology of Positive Economics* (1953), even if his positions, primarily in economic methodology literature, have been widely contended¹⁰. Without getting into details, theories are seen mainly as "filing systems" with which to arrange knowledge, and as instruments for making predictions. For the sake of the latter, one can build its models — Friedman says *must* — using hypotheses that are descriptively false. Models based on the assumption that individuals behave *as if* they were optimisers is an example.

The second view has not a comparable standard. Many impor-

⁹As in Paul David (1985). In two classical papers, Paul J. McNulty traced back to Cournot and Jevons the most important "basic conceptual change" that lays behind the two visions I sketched, namely, the passage from a view of competition as a behavioural activity to its conception as a market structure (1967, 1968). Adam Smith "did not conceive of competition as a 'situation' at all but, rather, as an active process leading to a certain predicted result. The Smithian concept of competition is essentially one of business behavior [...] The concept of competition originating with Cournot, on the other hand, is totally devoid of behavioural content. This is because Cournot's focus was entirely on the effects, rather than the actual workings, of competition" (McNulty 1967, 398).

¹⁰Hausman (1992) convincingly argued that mainstream methodological practice resembles more closely Mill's *a priori* inductive approach. Nevertheless, "although confused, mistaken, and inconsistent with the practice of mainstream economists" (Hausman, 1998), Friedman's views have been warmly embraced because of the "protective belt" they put around the core of neoclassical assumptions about individual's rationality, and around Robbins's identification of the "allocation of scarce means" as the predominant causal factor of interest in economic domain.

tant insights, however, can be found in Hayek's *Studies in philosophy, politics and economics*, mainly in *Degrees of Explanation* (1955) and *The Theory of Complex Phenomena* (1964).

His main point is that Economics methodology has to be clearly distinguished from that of the so called "hard" sciences, like Physics. About the latter, he agrees with Popper in saying that the direction of the explanatory path is from the "known to the unknown": Given a known phenomenon, the task of the scientist is to discover the unknown law by which it can be deduced and this way fully explained. This job, however, "presupposes conditions which are not present in many other fields" (Hayek 1955, 5). Disciplines like Economics, or Meteorology, force the scholar to a different path:

they are, in a sense, derivative: they consist of deductions derived from combinations of known laws of physics, and do not, strictly speaking, state distinct laws of their own but elaborate the laws of physics into explanatory patterns appropriate to the peculiar kind of phenomena to which they refer (*ivi*, 6-7).

we shall here have to proceed in our deductions, not from the hypothetical or unknown to the known and observable, but [...] from the familiar to the unknown (*ivi*, 9).

The reasons for which the familiar knowledge we have on individuals' behaviour ends up in an unknown macro outcome rests on the complexity of economic systems. The edge that switches the "reversal" of the explanatory path is certainly quite difficult to define. I would'nt dare to identify it — as Hayek does — in the fact that physicians study phenomena where the "number of significantly connected variables" is sufficiently small. In relative terms, however, the *comparatione* has sense. Murray Gell-Mann has been attributed a nice posing of the matter: "Imagine how hard physics would be if electrons could think"¹¹.

However we define it, it is the complexity of economic systems that calls for a different kind of investigation. We have a huge knowledge on the actors and institutions of which the economy is made. We know not to know many other "pieces" of knowledge, being them tacit, dispersed or hardly measurable. We know, finally, that even simpler systems where autonomous agents interact with

¹¹I due the anecdotic argument to Page (1999).

each other can easily give rise to highly complex dynamics. The "derivative" nature of the enquiry, then, has to be accompanied by the tentative selection of the relevant items of this knowledge. A reliable (scientific) explanation of the unwanted regularities that emerge at the macro level will then rest on

whether we have selected the appropriate hypotheses from our store of accepted statements and have combined them in the right manner (*ivi*, 7).

we do not invent new hypotheses or constructs but merely select them from what we know already about some of the elements of the phenomena; in consequence we [ask] whether the factors we have singled out are in fact present in the particular phenomena we want to explain, and whether they are relevant and sufficient to explain what we observe (*ivi*, 11).

We can stress two main features that according to Hayek characterise a compositive methodology. The first, is that the forecasting capabilities it gives are open-ended and miss details, since we can explain, and hence predict, only *kinds* of phenomena. The knowledge that one gets is knowledge of the *principle* of the things. Economics, in this view, suffers of the same limits of Darwin's evolutionary theory: One can explain fairly well the advent and survival of a certain organism, but as regards prediction he can point just to the kind of phenomena that will probably happen¹².

The second point worth noticing has to do with the "realism" of the building blocks of an explanation. Friedman's suggestion was that a theory "must be descriptively false in its assumptions" (Friedman 1953). Here, the precept is the opposite: One should not invent new hypotheses, but "merely select" them from the store of accepted knowledge about a phenomenon.

I will try to refine both conclusions later on — the precept of "realism", in particular, is somewhat a severe one. Anyway, the "core" of the arguments put forth by Hayek is sufficient to proceed, investigating whether ACE can play a role in this view of scientific explanation.

¹²Not to talk about long term forecasting: In the long term fishes can start to fly. It is of some interest how Hayek derives from these limits his personal view of policy making. "Knowledge of the principle of the thing will enable us to make circumstances more favourable to the kinds of events we desire [...] Such activities [...] should perhaps better be described by the term cultivation than by the familiar 'control'" (Hayek 1955, 19).

4 Was Hayek an ACE?

We now have an instance of why the traditional law-like view of scientific explanation is not inevitable. As Hayek put it, the "derivative" nature of the research, and the complexity of economic systems, have as a typical output something more qualitative: We have explanations of the principle, the reconstruction of the *explanatory patterns* that lead to certain kind of phenomena.

As regards the modelling strategy, this reconstruction will consist in the "selection" of some elements from the mess, and in the investigation of the result of their interaction.

Is ACE a suitable way of performing this kind of research? Some insights have recently been given by Vriend, in a paper where he addressed the question whether Hayek could be considered an ACE *avant-la-lettre* (1999). As for as the topics investigated, he points to the many themes typical of the Austrian economist present in ACE literature — from his view of the economy as a complex evolving system, to the stress on the narrow limits of human reason, on rule-based behaviours, and on the processes of knowledge diffusion. As regards the approach, he argues that ACE research can be considered an application of Hayek's methodological insights.

One of his main stresses is in the fact that simulations *are* models with which to represent economic systems, in that they are abstraction from reality. To them, hence, can be applied Hayek's insights on the status of models in economic reasoning. Mainly, the fact that they are hardly verifiable, if not by means of consistency checks between the models' outcomes and the phenomena under investigation:

social theory is explanation, and explanation means modeling. Models can be presented in various forms. They could be either purely verbal or quantitative, where quantitative models, in turn, might be a set of mathematical expressions or a set of computer instructions. That is, a computer program as used in ACE *constitutes* a model. And since explanation *is* modeling, and this is what social theory is all about, an ACE computer program as such *is* social theory. The only, and essential, reason to execute the computer program is to carry out the consistency checks. (Vriend 1999, 31-32).

The point is not trivial. Computer languages are not "outliers" in scientific reasoning: They are just one of the symbol systems

available to the social scientist (Gilbert and Terna 2000, Ostrom 1988). It should be noted, however, that to clarify that computer languages *are* among the symbolic systems available to build models is not sufficient to assess the methodological relevance of that kind of models. Also paintings are abstractions of reality, but they are hardly relevant models for the economist — while it plainly could be the case to the historian, think of *Guernica*.

More precisely, it is true that explaining implies some sort of modelling, and that simulations are a suitable language to do this job. But for a model to become a relevant explanation we need some more steps: We need a "template" of what we bear for a scientific explanation, and the check that our model fits with the template. A clarification of this point, together with the arguments given in previous section, will be sufficient to reinforce Vriend's parallel between ACE models and Hayek insights.

Consider, as an example, whether a certain model can be used to build a proposition A coherent with the "covering law" view of scientific explanation. The first check is that A , and hence the model encapsulated in A , has to be written in an appropriate language. The second check is about syntax, and requires that A has the following template: " $L \cap P \Rightarrow E$ ", where L (the model) states a general law, P is a proposition asserting the presence of particular events, and E is the *explanandum* — the proposition stating the fact that has to be explained. The last check has to do with semantic, and requires for A to be passible of an empirical test. The check, under this view, requires that P and E were propositions with an empirical content, i.e., of which we can ascertain whether they are "true" or "false". If all the checks have been passed, then we can say that A is a (in principle falsifiable) scientific explanation.

As regards the relevance question about a simulative model, then, once we have recognised that it has been written in a suitable language, there are two more points to be addressed. The first, is the formal check: Is our model an instance of what we bear for a "scientific explanation"? The important contribution of Hayek, that I synthesised in previous step, is that when addressing a certain kink of phenomena, a suitable explanation may consist in 'putting together' something that we know about the phenomenon and see which kind of aggregate behaviour emerge. And simulations can do that job.

This way, we produce the following hypothesis: The *structure of relations* that we built — and *not* the quantitative connection between micro and macrobehaviour — is akin to the principle at work in the real world.

The second point to be addressed, then, is how to test this hypothesis¹³. The hints we can derive from Hayek point to something by necessity qualitative. *If* the aggregate behaviour emerging by our model is (not) consistent with what we observe in reality, *then* we will have an evidence corroborating (disconfirming) the relevance of the factors we selected.

It should be clear then, that agent-based simulations with a sufficiently clean design are certainly a way of implementing such a compositive kind of investigation¹⁴, and that under this view they *are* well formed scientific explanation of the phenomena studied.

The power of the explanatory pattern hypothesised will not rest in the statement of a "well-behaved" law with which to deduce a general class of occurrences, but in the anecdotic evidence that the certain set of assumptions that has been singled out is *sufficient* to generate the emerging behaviour. This is not a direct confirmation that we grasped the "true" causal mechanisms, but if we have been able to mimic the regularities observed, says Hayek, "there is good reason to regard the model as exhibiting the principle at work in the more complex phenomenon".

Needless to say, simulations are just the last-comers among the possible instantiation of the methodological insights of Hayek; verbal reasoning such as a mathematically formalised one are sound competitors. The relation between ACE and Hayekian views, moreover, is not that of an inclusion. To proceed in our investigation, then, we have to relax the link we traced with Hayek's insights and generalise some of the hints given.

¹³Vriend reports a rather pessimistic view of Hayek on this matter: "a simple theory of phenomena which are in their nature complex [...] is probably of necessity false" (Vriend 1999, p. 30). The quote is from *The theory of complex phenomena*.

¹⁴The clear design is again a crucial requirement. As Kirman and Vriend put it, "when modeling by building artificial worlds, one might get a very good fit without gaining understanding. There exist economic simulation models with more than 10,000 variables. At some point it might be that one mainly succeeds in building a copy of the real world, about which we have the same degree of understanding as about the real world.", mimeo, quoted in Vriend 1999).

5 The methodological status of agent-based simulations

In the arguments given so far I left (at least) three points somewhat sketched. The first relates to the conception of scientific enquiry as a mean to drive predictions. What about simulations? We saw that, as for evolutionary theory, the link between a suitable explanation and prediction can be weak. We saw also that, this notwithstanding, it is conceivable to turn an explanation of the principle into policy indications in a "cultivation" style¹⁵. The point to be addressed here is more general, and regards what can be called the pragmatic content of theories.

Secondly, it is not obvious which is the logical structure (the template) of what Hayek called "explanations of the principle". More in general, a clarification of the logical structure of a simulative approach is essential to assess the relations between this kind of theorising and empirical tests.

Finally, I based my arguments on a partial reconstruction of Hayek's methodological positions, and this way a key role has been played by the realism of assumptions, by their selection from a "store of accepted statements". To a certain extent, this might not be a limit: Strictly speaking, in that store we could even find some *as if* assumptions, provided that, say, the maximising behaviour of individuals could be considered an accepted conception. The subjectivist approach that can be recognised in all Hayekian research, however, stays at odds with this interpretation. To impose the realism of assumptions, therefore, can be demanding; moreover, it does not match with ACE practice: A growing strand of ACE literature is focusing on the dynamics emerging in artificial markets populated with zero-intelligence agents. If we agree that this is not a realistic assumption, a brief discussion on the theme is needed.

In this section, I will first tackle with these points discussing two topics essential for my goals: That of abductive inference, and that of the realism of assumptions. In the last paragraph I will cast the arguments given into some results of modern theory of abductive inference, providing this way a framework that helps clarify the logical status of simulations and their role as causal explanation of economic phenomena.

¹⁵See footnote 12.

5.1 Abductive inference and the pragmatic content of theories

To address the first two points, it is useful a brief digression on a logical and philosophical background that characterise agent-based modelling, and that is seldom recognised. Curiously enough, one of the "main enemies" of ACE partisans, Friedman's *as-if* construal, can be traced back to roughly the same philosophical roots. Let us start from this latter.

In the received view, the so-called "symmetry principle" between explanation and prediction holds, since they both refer to the same syllogism " $L \cap P \Rightarrow E$ ": Given E , its *explanation* consists in showing its link with a law L and the particular conditions P ; if we bear a law L for true, on the other hand, the occurrence of P allows us to *predict* that E will happen. Where this principle were true, we had a clear link between positive economics and policy indications; in Auguste Comte words, "Science hence prediction, prediction hence action."

As a matter of fact, the stress on the forecasting role of theories typical of the approach put forth by Friedman is all but marginal. What about the roots of this approach? The first comprehensive attempt to clarify the underlying philosophy of *The Methodology of Positive Economics* is due to Boland (1979,). He recognised in Friedman's view the main tenets of instrumentalism, a term coined by John Dewey to tell his methodological lecture of Charles Sander Peirce's pragmatism from the metaphysical one given by philosophers like William James or Ferdinand Schiller¹⁶.

The key feature of this approach is well clarified by Popper: "By instrumentalism I mean the doctrine that a scientific theory [...] should be interpreted as an instrument, *and nothing but an instrument*, for the deduction of predictions of future events" (1983, 237)¹⁷.

¹⁶In turn, to tell his views from the both, Peirce coined the term "pragmatism". See Rossi (1986, 234-235), Apel (1981), Abbagnano (1968, *ad vocem*). On the attribution of these roots to Friedman's positions agrees also Mark Blaug: "in fact he owes more to John Dewey than to Karl Popper: he is a pragmatist rather than a falsificationist" (1992, 104). Interestingly enough, however, the *as-if* argument has originally been proposed by a philosopher of the "metaphysical" wing, Hans Vaihinger, in his *Die Philosophie des Als Ob (The philosophy of as-if, 1911)*.

¹⁷Compare also Feyerabend: "the view that scientific theories are instruments of prediction which do not possess any descriptive meaning" (1981, 144). Roughly the same definition is used by Boland in its arguments.

Peirce's pragmatism, however, was less demanding. Simplifying a bit, the idea that theories have an "instrumental" value stems from his view that the meaning of any conception consists in its *use*. In this broader sense, to limit the pragmatic content of a theory to the "deduction of predictions of future events" is all but inevitable, under two respects.

The first one has to do with the "use" of a theory as an explanatory tool. *If* the symmetry principle hold, there would be no questions. To the extent that an explanation should intelligibly answer to a question "why", however, there may be a trade off between explanatory power and forecasting capabilities. Consider for instance the following theory L : "Macroeconomic time series follow an AR(1) data generating process". Given a certain value of a macroeconomic variable E , its explanation will amount to compute the "particular conditions" P , in the form of the parameters of the autoregressive process that best fit the time series of that macroeconomic variable. With this theory at hand, predictions of future events are possible, but the question "Why is this happening?" has not been answered. On the opposite, it is also conceivable a purely idiosyncratic explanation, in the form of the historical reconstruction of the causal chain that produced E . Here, we still (implicitly) use some "laws", with a form like "in such circumstances, these facts usually occurs"; but the possibilities of drawing predictions are particularly weak. To the extent that we recognise a *per se* value to an explanation, we *use* theories to clarify phenomena whichever their forecasting power.

About the use of our knowledge by means of *deductive* syllogisms, moreover, Peirce argued that this kind of inference is the less frequent in ordinary reasoning, and just one of the possibilities in scientific investigation. Prior to Peirce contributions, logicians commonly divided arguments into two classes: Deductive arguments (a.k.a. necessary inferences), and inductive arguments (a.k.a. probable inference). A most important tenet of Pierce's pragmatism resides in the stress he put in another form of probable inference, that he called abduction (or retroduction, or hypothesis). Technically, abduction is the inference of the minor premise P from the *explanandum* and the general law: " $L \cap E \rightarrow P$ ".

A clear way to look at abduction is in sampling theory. Let the three items of a syllogism be the following:

$L \equiv$ "All beans in urn A are white";

$P \equiv$ "The beans of sample s have been drawn from A ";

$E \equiv$ "The beans of sample s are white".

Deduction is the (necessary) inference of E by L and P . Induction is the (probable) inference of the rule L by E and P . Abduction, is the (probable) inference from L and E to the case P . Knowing that the beans of sample s are white, and that all beans in urn A are white, we hypothesise (abduce) that s has been drawn from A .

In the years, he extended this early characterisation of inference forms into a broader view of the scientific method, which main features are the three steps abduction-deduction-induction, together with what he called the "economics of research" (Burch 2001). Abduction became defined as inference to, and provisional acceptance of, an explanatory hypothesis; that is, inference to something that clarifies or makes routine some information that has previously been surprising: "The surprising fact, E , is observed, but if P were true, E would be a matter of course, hence there is reason to suspect that P is true"¹⁸. The provisional nature of this belief stems from the fact that something else could explain E , and that the causal link between P and E might not be necessary.

In the Seventies, Harry Pople (1973) pointed out the importance of abduction to the literature on artificial intelligence, and from then on it has been applied to fields like medical diagnosis, natural language understanding, plan recognition, and many others (McIlraith 1998; Menzies and Gambetta 1994). To the economists, probably the most interesting case is that of signals extraction: The diffusion of knowledge *via* price signals is exactly the inference of an information (for instance, the rise in the profitability of a firm) from the observation of the typical *effects* of the information abducted (in the example, the rise in the firm's stock price). A rational investor has no way to *deduce* that a rise in the profitability of the firm has occurred. But, since a rise in the stock price is a typical effect of a firm's profit chances increase, it can be used to hypothesise that its profitability has actually risen.

¹⁸Peirce (1955); I changed the lettering for ease of readability. The term "suspect" has a particular fit: Probably the most famous example of the abductive inference of explanations is that of Conan Doyle's Sherlock Holmes.

These examples can be relevant when we model an agent behaviour. The point to be stressed here, however, stand somewhat at a higher level. Namely, abduction is a key step also in scientific theory formation and validation (Thagard 1988). When we ask *why* a certain phenomenon occurs, we will typically proceed 'putting together' a constellation of possible causal factors, trying to reproduce the phenomenon observed. The acceptance of a hypothesis, the theory validation, is based on how well it reproduces the evidence, and on other criterions — such as the simplicity, or some measure of plausibility of the explanation — with which to compare competing explanations. A modern synonym for abduction, i.e. inference to the best explanation, well reflects this research path.

Concluding, the stress on the "deduction of predictions of future events" germane to the instrumentalist *as-if* construal has certainly sound bases: It stems from the need, "built-in" in scientific reasoning, of producing conceptions with a pragmatic content. However, the pragmatic content of scientific knowledge can typically resides elsewhere, namely, in its use *via* abductive inference to build explanatory patterns that may or may not have an immediate "surplus" of forecasting power.

We can turn now to discuss which kind of hypotheses we can buy in our "store of accepted statements".

5.2 The realism of assumptions

Its old age notwithstanding, the quarrel about the realism of assumptions has not yet arrived to a sharp solution. The many tricks that it hides can be clarified starting with the following distinction proposed by Machlup:

one must distinguish two kinds of protest against lack of realism: first, that certain realistic assumptions were omitted [...]; second, that the assumptions made were unrealistic.¹⁹

Let's say that Hayek would endorse at most the first form of "unrealism", while Friedmann is a supporter of the second one²⁰. Hayek tells us: You certainly have to ignore some factors present in the system you are studying, but the other building blocks of

¹⁹Machlup 1978, p. 78.

²⁰Needless to say, the opinions of both authors that I will report are a raw proxy of the originals.

your theory must conform with what you know about that system's elements. Friedmann tells us: You do not have to look into the black box, the only thing that matters are the outcomes of your theory.

Let us try to follow Hayek's suggestion, for instance, while building some model of unemployment. Let us omit the explicit consideration of, say, money and politicians: The only actors of our model will be firms and households. The approach we are adhering to, impose us to make realistic assumptions about them. The point is: What is a realistic assumption, for instance, about a household? If we are not in an experimental economics department, we will have to use a model of it. To build a realistic model of a household, our approach allows us to omit some relevant factors that affect households' behaviour, but all the other assumptions that we make about them have to be realistic. Let us say that we omit the explicit consideration of gender differences, and that we model our family simply as the union of a variable number of individuals with a unique, common budget. Our rule says that we are allowed to do this, but the individuals we use have to be realistic. And since we are not in an experimental economics department, we are obliged to postpone again our "realistic check" and deliver it on the model we will choose to represent them. And the game starts again. The point is, that we always have to code in some way the objects we are studying, and hence the problem of "realism" will recursively reappear inside any assumption we make. The best place where we can end up, following blindly this approach, is in a gloomy reductionism: the explanation of everything is always in some deeper layer.

Have we to endorse the "black box" approach to eventually stop this loop? It is not that simple. A widely recognised stylised fact about economists' efforts to account for that share of active population that has not a job, is that there are many competing theories which do possess some fit on the data. To impose that each of these theories is a black box where it is forbidden to look into seems a bit punishing, since some scrutiny of the assumptions could at least in principle helps us in the choice among the competitors. To strictly adhere to the black box approach, in other terms, could easily leads us in another gloomy place, where the explanation of everything can be very smart, but is somewhere you can not look at.

If it were possible to investigate any subject in an experimental manner, or if it were always possible to choose between competing

theories thanks to clear answers from a dataset, one of the two approaches would be suitable. As it seems, that is not the case.

As Hausman put it, we have not yet solved Mill's problem (1992, 1998). From one side, Economics has characteristics of "separateness" that justify the adoption of idealised assumptions, unrealistic in some strong sense. On the other side, any pure theory conflicts with the "inexact" nature of real world economic relations. Given the difficult interplay between theory and empirical evidence, a scrutiny of the single assumptions adopted seems unavoidable²¹. When we do not overstate the scope of applicability of economic theorems, however, Hausman's restatement of Mill's problem provides a useful clarification of their limits: In the case we recognise a certain degree of conformity of our assumptions to what is there in the real world, an economic law will tell us which are the *tendencies* that will be at work.

A similar position has recently been endorsed by Robert Sugden (2000). He analysed some "high level" features of two papers that, as he says, nearly all economists would be glad to have in their list of publications: Akerlof's *The Market for 'Lemons'*, and Schelling's *Dynamic Models of Segregation*. Both works share an apparent contradiction: Their authors point to some real world facts that are very closed to what they study, but nevertheless they do not pretend to explain any empirical evidence. They present their work as a highly idealised one, without using the highly idealised language we would expect in the presentation of a pure formal system. The contradiction vanishes when considering that both works pointed out the potential effects of causal factors that can rarely be observed in isolation, but nonetheless can be essential to fully comprehend real world situation in which the tendencies they identify are at work.

If we consider sufficiently fair this vision of what economists really do when they put together various assumptions, we should recognise

²¹Mill's own solution was to consider what we usually call "first principles" as "last principles", hence considering the scrutiny of assumptions a sort of fundamental step that give them the status, with respect to the rest of economic theory, of *a priori* truth. "Instead of being the fixed point from whence the chain of proof which supports all the rest of the science hang suspended, they are themselves the remotest link of the chain. Though presented as if all other truths were to be deduced from them, they are the truths which are last arrived at; the result of the last stage of generalization, or of the last and subtlest process of analysis, to which the particular truths of the science can be subjected; those particular truths having been previously been ascertained by the evidence proper to their own nature." (1974, 88).

to simulations a better equilibrium between pros and cons than that I used at the beginning of the paper.

In particular, a computer simulation *can* be considered a way of delivering experimental evidence, since — typically — economists are not interested only in the behaviour of "real" individuals, but also in the effects of some idealised behaviour. In sociologists terms, we study roles, not persons, and roles can fairly be coded into rules and other formal stuff: Actually, it is what most economists do. By the same token, they can be coded into "pieces of software", and in this way they can be observed during an experiment about some topic of interest.

It is not in the capability of simulations of managing realistic agents or institutions that resides their first added value, but in the fact that they can bring the more or less realistic assumptions made by a scholar into a laboratory to test in a "controlled way" their outcomes. This laboratory is a "true" laboratory, in the limits that we consider economics as a separate discipline handling with the effects of just some idealised aspects of economic behaviour.

5.3 The simulative route to causal explanations

I started this paper with a sharp contrast between a law-chasing approach to economic reasoning and the anecdotic outcomes of agent based simulations. I proceeded showing a possible point of contact among the two approaches, and discussing other topics relevant to the clarification of their logical status. To close the argument, I will now proceed showing to which kind of questions both ACE and traditional modelling are suitable methodologies to answer to. Once the two approaches are this closer, it will be possible to appreciate also their different features.

As regards law chasing, I pointed out that when the concern is on market process more than on equilibrium outcomes, and in the aggregation of individual behaviours more than on optimal allocation problems, we will typically end up with something more fuzzy than general laws. As a mnemonic, call this kind of questions "fuzzy questions". The fact that our output can not be a precise law does not imply that algebra is not a suitable language with which to explore these matters. It simply means that the "laws" that we find

out have to be considered like qualitative statements about structure of relations, without a precise quantitative content. Both Hayek and Keynes already pointed out this fact. According to the former, the knowledge that resides in any model about complex phenomena rests in the "certain range of phenomena which can be produced by the type of situation which it represents. [Such a model] will contain assertions about a *structure of relations* even if we do not know the value of any of the variables."²² Theories, then, are "destined to remain 'algebraic', because we are in fact unable to substitute particular values for the variables, [and cease] then to be a mere tool and becomes the final result of our theoretical efforts."²³ In a letter to Harrod about the role of econometrics, Keynes expressed roughly the same concept:

In chemistry and physics and other natural sciences the object of the experiment is to fill in the actual values of the various quantities and factors appearing in an equation or a formula; and the work when done is once and for all. In economics that is not the case, and to convert a model into a quantitative formula is to destroy its usefulness as an instrument of thought.²⁴

The arguments of Hausman and Sugden that I reported in last section point in the same direction: At the end, once we abandon the pretention of arriving to golden rules everywhere and every time applicable, theorem building can be viewed simply as an analytical engine with which to build causal links between a certain stylised aspect of economic systems and some aggregate outcomes. In a sense, the generality of an economic theorem is just a byproduct of the peculiar tool we used to recognise a causal relation.

As a matter of fact, one of the early critiques that have been raised to the covering law approach, by authors like David Scriven or Wesley Salmon, is that it missed to recognise to the concept of causality its key role in scientific explanation²⁵. The shift of focus has important consequences, since when we put the causality at the centre of the stage the requisite of providing universal arguments is no more an essential one.

²²Hayek 1955, p. 15, emphasis added.

²³Hayek 1964, p. 28.

²⁴Reprinted in Hausman 1994 p. 301.

²⁵See Scriven (1962, 1963) and Salmon (1984). A wide synthesis on the history of the concept of causality can be found in Laudisa (1999). For the relations between causality and the D-N model see Salmon (1990).

As a first approximation, we can say that for a certain factor P to be recognised as the cause of an *explanandum* E we simply need to show that P is a sufficient condition for E to happen. The sufficiency of P , then, can be recognised by means of a formal analysis, but can also be ascertained *in machina*: If P is sufficient to generate E as an outcome, this will be the case also in a simulation. The anecdotic nature of the evidence produced, indeed, does not infringe the causal link recognised, since for the latter to be verified we need just the experiment to be replicable.

Actually, it is possible to conceive of causality also in terms of a necessary relation, and in this case some law like statement would again be essential. John L. Mackie (1962), proposed an explication of the concept capable of clarifying the role both of the necessary and the sufficient conditions. He considered the case of a fire in a building, of which a short circuit is recognised as the cause. In general, a short circuit can not be considered a necessary condition for a fire, since many other factors can start it. Neither it can be considered a sufficient one, since it can happen, say, far from flammable materials. How should we interpret, then, the proposition "The fire has been caused by a short circuit?". Mackie proposed the following explication: First, even if a short circuit is not a sufficient condition for a fire, it came with a set S of events that, as a whole, is a sufficient (even if not necessary) condition for a fire. Second, the short circuit is an essential element (a necessary one) of the set S , meaning that, without it, the set S will no more be a sufficient condition for a fire. With somewhat a short circuit of words, he proposed to consider the short circuit an Insufficient condition, Necessary to an Unnecessary but Sufficient condition for a fire to happen; in a word, an Inus condition.

Typically, even if we are unable to provide definitive laws about fuzzy questions, we can give them provisional answers looking for some Inus conditions. Consider the question "What causes business fluctuations?". A possible answer, is to show analytically that random shocks, together with some adjustment mechanisms, are a sufficient condition for business fluctuations to arise. Neither the random shock alone, nor the deterministic part of the model, are sufficient to generate fluctuations, and on the other side cycles can emerge also as a consequence of other causes. With a simulation one can do pretty the same job. With a lower degree of generality,

they nevertheless provide a suitable explanation for the phenomenon observed.

What are, then, the differences? In what agent based simulations can provide additional insights? In order to answer, it is useful to recall the Peircean notion of abduction, providing a modern formalisation of it. Following Bylander et al. (1991), I will use the following notation. d stands for a datum, e.g., a fact to be explained, and D stands for a set of data. h stands for an individual hypothesis, e.g., a causal factor, and H stands for a set of hypotheses, or a composite hypothesis. We can now give the following formalisation of abduction, or of inference to a best explanation.

An *abduction problem* is a tuple $\langle D_{all}, H_{all}, e, pl \rangle$, where:

D_{all} is a finite set of all the data to be explained;

H_{all} is a finite set of all the individual hypotheses;

e is a map from subsets of H_{all} to subsets of D_{all} (H_{all} explains $e(H_{all})$);

pl is a map from subsets of H_{all} to a partially ordered set;

pl can be viewed as a measure of plausibility, of belief, or of relevance given to the causal factor. We will say that:

H is *complete* if it explains all the data, i.e. if $e(H) = D_{all}$;

H is *parsimonious* if no proper subset of H explains all the data that H does;

H is an *explanation* if it is complete and parsimonious;

H is a *best explanation* if no other explanation is more plausible than H , i.e., it does not exist a $H' \neq H : pl(H') > pl(H)$. It is only "a best" because of the partial order imposed to H : Hence, there can be several best explanations.

A clear way to think of this kind of inference is in medical diagnosis. The sets D are collections of symptoms, while the set H_{all} is the collection of all the causal factors that produce some symptom in the organism. We can think of scientific research as the activity of filling up the set H_{all} and keeping up to date the map e , with the

list of symptoms produced by the new causal factors we are adding. The diagnosis, is the use of this knowledge to infer an explanation to a certain set D_{all} of symptoms observed in a patient.

In economics, such a clear distinction is not possible, since we unfortunately have not a "neutral" organism linking a causal factor in isolation with its outcomes; say, a certain production function with the level of output. For the sake of simplicity, however, let us assume that we already have a received view on some basic features of the "economic organism", and on its "normal" functioning. We can think, for instance, that a "healthy" economy is one with firms and individuals perfectly rational and informed, with given technology and tastes, that is developing along an optimal growth path. The "diseases" d range from different technologies and tastes, exogenous dynamics for some relevant variables, market imperfections, information asymmetries, anomalous (not rational) behaviours and so on. For each of these, we are interested in which are the "symptoms" they produce, i.e., the modifications to the normal behaviour of our healthy economy.

Here, we can find a first difference between theorem building and simulations. If we are smart enough, with some algebra we can ascertain the effects of all the elements of a (large) subset H at a time. For instance, the effects on the benchmark behaviour of an entire class of different technologies. If the algebra does not help us — and we are smart enough — we can turn to simulations and ascertain the effects of some additional factors, but just few at a time. The difference is: Algebra is the coupée, and simulations are the jeep. Theorem building has a high productivity in parsing the territory H_{all} , simulations allow us to go beyond and investigate some additional h .

Consider now the other task, the "diagnosis". We observed a phenomenon D , and we want a best explanation to it. In this case, we have to explore our set H_{all} looking for a condition h sufficient to give rise to D . The first thing to say is that, if we do not find a single $h : e(h) = D$, it is probable that we are already facing a fuzzy question. In fact, as a next step we must try various h together, hoping to find a composite hypothesis sufficient to give rise to D . The point is, that typically there is some non trivial interaction between various h , where "trivial" means that the factors considered are not mutually exclusive, and that their effects are additive. If

the interaction among the h is not trivial, Bylander et al. (1991) provide an interesting complexity result. They show that in general the exploration of the set H_{all} becomes an NP -hard procedure: In other words, the search for a best explanation is a not computable task.

How is it that we usually keep many economic models for "good" explanations, then? A provisional answer, rests on the hypotheses we make about the set H_{all} and its relations with what I called a benchmark, "healthy" economy. In principle, economic theory has as outcomes not only the mapping between the sets H and D , but also many restrictions on the plausible hypotheses h . That is to say, that the set H_{all} that we have to parse can be reasonably small, and hence, even if the search for a best explanation is potentially long, its effective duration is limited by the small number of alternative hypotheses that we are willing to consider. On the other hand, however, the benchmark to which I referred is not a datum, since we are not interested only in the effects of single causal factors, but, so to speak, in the mechanism linking this factor to certain outcomes.

To turn the formalisation adopted into a fair portrait of how works scientific explanation in economics, then, we can say the following. The hypotheses we adopt about the mechanism — in the metaphor, about the normal working of the economic "organism" — is a matter of research programme selection. That is to say, we choose a hard core of hypotheses H^* that works as a basic framework, and do "normal science" studying the effects of a fairly limited number of additional hypotheses and initial conditions. The partial ordering pl , hence, is defined only among the sets H that includes H^* . Due to the limited number of the elements among which an order relation is defined, a best explanation, provided that an explanation at all exists, can in principle be found.

Here is the second difference between traditional modelling and agent based simulations. The clause "for the sake of tractability" has, as a by product, the limitation of the scope of composite hypotheses that one can test. If we add to this the adoption of a reference framework (a hard core) of hypotheses about the main features of the system studied, the inference to a best explanation about a certain phenomenon is a well defined task; somehow, the potential complexity/uncomputability of spanning a wider domain is

avoided by definition. Needless to say, however, it is not a mechanical task: An innovative theory is often based on the identification of a new relevant factor h more than on the mechanical recombination of what is at the stake in the literature.

With an agent based simulation, the high level target is the same: To find out a set H that is an Inus condition for the emergency of a pattern D . The logical structure of the explanation provided, hence, can fairly be considered the same. The task, however, is a more complex one under two respects. On one side, we are driving a jeep, and the span of alternative hypotheses is slower. The clause "for the sake of tractability", moreover, is no more active, and the search for a best explanation can easily end up in a NP -hard intractability.

From a methodological point of view, then, agent based models can have a potentially low explanatory power not because of a lower degree of generality, but because of the higher degrees of freedom they grant in the combination of elementary causal factors. This freedom, ends up in a more difficult comparison with alternative patterns of explanation.

An essential step in fostering an agent based approach, hence, is the adoption of a plausibility measure pl alleviating the computational requirements necessary to the exploration of the set H_{all} . Not surprisingly, such a rule has spontaneously evolved in the literature. From Axelrod's simplicity principle to the literature on zero-intelligence agents, one of the main tenets of ACE is the effort to give rise to macrobehaviours with "minimal" requirements to the rationality of the agents and to the architecture of the simulations. Epstein (1999) stressed this point stating a sort of first principle to a research programme: "The point is not how much rationality there is at the microlevel, but how little is enough to generate the macroequilibrium".

This sort of commitment to simplicity can play a role similar to the "for the sake of tractability" clause: It reduces the number of potential patterns of explanations to be considered. Hence, it defines both a research programme and some feasible limits within which a suitable explanation can be looked for.

Even if the simplicity of a simulation has this key role, the degrees of freedom a scholar has in spanning the possible interactions among different causal factors h -s are obvious.

5.4 Concluding remarks

I started with a sharp contrast, and ended with two similar worlds. The points of similarity between a "traditional" approach and a simulative one can be synthesised in the following: Both are concerned with some idealised portrait of economic agents/institutions; both are a way to discover causal linkages between this idealised individual behaviour and aggregate outcomes; both face the same potentially uncomputable task of finding a "best explanation", according to some plausibility measure; both overcome this complexity self defining the limits within which to test (to experiment) additional hypothesis.

The main differences, from a methodological point of view, rest in the way the first two points are handled. The cause chasing, within an algebraic approach, typically ends up with a system of equations defining the structure of relations that grasps — even in a simultaneous way — the causal linkages between the different elements. The "soundness" of these linkages is granted by the generality of the results obtained. Within a simulative approach, the structure of relations is built from the bottom up, in a more explicit way. The soundness of the causal links pointed out rests on the replicability of the experiment.

From a practical point of view, the differences are the "usual" ones. The trade off is between general results in an elegant packaging, and experimental evidence, sometimes more difficult to interpret, but much wider in the scope it gives to the research.

References

- [1] Nicola Abbagnano. *Dizionario di filosofia*. Utet, Torino, 1968.
- [2] Dario Antiseri. *Trattato di metodologia delle scienze sociali*. Utet, Torino, 1997.
- [3] Karl Otto Apel. *Charles S. Peirce: From Pragmatism to Pragmaticism*. University of Massachusetts Press, 1981.
- [4] Robert Axelrod. *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton University Press, Princeton, 1997.
- [5] Mark Blaug. *The methodology of economics. Or how economists explain*. Cambridge University Press, Cambridge, 1992.

- [6] L.A. Boland. A critique of friedman's critics. *Journal of Economic Literature*, 17(2):503–22, 1979.
- [7] Robert W. Burch. Charles sander peirce. *Stanford Encyclopedia of Philosophy*, page ad vocem, 2001.
- [8] Tom Bylander, Dean Allemang, Michael C. Tanner, and John R. Josephson. The complexity of abduction. *Artificial Intelligence*, (49):25–60, 1991.
- [9] Paul David. Clio and the economics of qwerty. *American Economic Review*, 75:332–7, 1985.
- [10] Joshua M. Epstein. Agent based models and generative social science. *Complexity*, (May):55–77, 1999.
- [11] Joshua M. Epstein and Robert Axtell. *Growing Artificial Societies - Social Science from the Bottom Up*. MIT Press, Cambridge (MA), 1996.
- [12] Paul Feyerabend. *Realism, rationalism and scientific method: Philosophical papers*. Cambridge University Press, Cambridge, 1981.
- [13] Milton Friedman. *Essays in positive economics*. University of Chicago Press, Chicago, 1953.
- [14] Milton Friedman. The methodology of positive economics. In *Essays in positive economics*. University of Chicago Press, Chicago, 1953.
- [15] Neil Gilbert and Pietro Terna. How to build and use agent-based models in social science. *Mind & Society*, (1), 2000.
- [16] Daniel M. Hausman. *The Inexact and Separate Science of Economics*. Cambridge University Press, Cambridge, 1992.
- [17] Daniel M. Hausman, editor. *The Philosophy of Economics: An Antology*. Cambridge University Press, New York, 1994.
- [18] Daniel M. Hausman. Philosophy of economics. In Edward Craig, editor, *Routledge Encyclopedia of Philosophy*. Routledge & Kegan Paul, London, 1998.
- [19] Friederich A. Hayek. Degrees of explanation. 1955.
- [20] Friederich A. Hayek. The theory of complex phenomena. In *Studies in philosophy, politics and economics*. Routledge & Kegan Paul, London, 1964.
- [21] Carl Gustav Hempel. *Aspects of Scientific Explanation*. Free Press, London, 1965.
- [22] Carl Gustav Hempel. *Philosophy of Natural Sciences*. Prentice Hall, New Jersey, 1966.

- [23] Carl Gustav Hempel and P. Oppenheim. Studies in the logic of explanation. *Philosophy of Science*, 15:135–75, 1948.
- [24] Kenneth Judd. Computational economics and economic theory: Substitutes or complements? (208), 1997.
- [25] Fritz Machlup. *Methodology of Economics and Other Social Sciences*. Academic Press, New York, 1978.
- [26] John L. Mackie. Causes and conditions. In E. Sosa and M. Tooley, editors, *Causation*. Oxford University Press, Oxford, 1962.
- [27] Sheila A. McIlraith. Logic-based abductive inference. Technical Report KSL 98, Stanford University - Knowledge Systems Laboratory, 1998.
- [28] Paul J. McNulty. A note on the history of perfect competition. *Journal of Political Economy*, (4):395–9, 1967.
- [29] Paul J. McNulty. Economic theory and the meaning of competition. *Quarterly Journal of Economics*, 82(4):639–56, 1968.
- [30] Tim Menzies and Windy Gambetta. Exhaustive abduction: A practical model validation tool. In ECAI, editor, *Workshop on Validation of Knowledge-Based Systems*, 1994.
- [31] T. Ostrom. Computer simulation: the third symbol system. *Journal of Experimental Social Psychology*, (24):381–392, 1988.
- [32] Scott Page. Computational models from a to z. In Santa Fe Institute, editor, *SwarmFest 1999*, 1999.
- [33] Charles Sander Peirce. Abduction and induction. In J. Buchler, editor, *Philosophical Writings of Peirce*. Dover, New York, 1955.
- [34] Harry E. Pople. On the mechanization of abductive logic. In , editor, *Proceedings of the Third International Joint Conference on Artificial Intelligence*, Stanford, 1973.
- [35] Karl R. Popper. *Poscritto alla Logica della scoperta scientifica. 1. Il realismo e lo scopo della scienza*. Il Saggiatore, Milano, 1983.
- [36] Paolo Rossi. *Le filosofie del novecento*. Sansoni, Firenze, 1978.
- [37] Wesley C. Salmon. *Scientific Explanation and the Causal Structure of the World*. Princeton University Press, Princeton, 1984.
- [38] Wesley C. Salmon. *Four Decades of Scientific Explanation*. University of Minnesota Press, Minneapolis, 1990.
- [39] M. Scriven. Explanations, predictions, laws. *Minnesota Studies in the Philosophy of Science*, 3:170–230, 1962.

- [40] M. Scriven. New issues in the logic of explanation. In S. Hook, editor, *Philosophy and History*. New York University Press, New York, 1963.
- [41] Robert Sugden. Credible worlds: The status of theoretical models in economics. *Journal of Economic Methodology*, 7(1):1–31, 2000.
- [42] Pietro Terna and Rosaria Conte. Una discussione sulla simulazione in campo sociale: mente e società. *Sistemi Intelligenti*, 12(2), 2000.
- [43] Leigh Tesfatsion. Agent-based computational economics: A brief guide to the literature. <http://www.econ.iastate.edu/tesfatsi/>, 2000.
- [44] P. Thagard. *Computational Philosophy of Science*. MIT Press, Cambridge (MA), 1988.
- [45] Hans Vaihinger. Die philosophie des als ob. unpublished, 1911.
- [46] Nicolaas J. Vriend. Was hayek an ace? mimeo, 1999.