Causality, emergence, computation and unreasonable expectations

Fabio Boschetti, CSIRO

Abstract

I argue that much of current concern with the role of causality and strong emergence in natural processes is based upon an unreasonable expectation placed on our ability to formalize scientific knowledge. In most disciplines our formalization ability is an expectation rather than a scientific result. This calls for an empirical approach to the study of causation and emergence. Finally, I suggest that for advances in complexity research to occur, attention needs to be paid to understanding what role computation plays in this experimental approach.

1 A thought experiment

Let's suppose we want to understand social behavior as a function of the interaction among individuals and we plan to use a computer model in our study. For example, we may want to test the hypothesis that current global market behavior is the result of economic agents acting "rationally," as suggested by classic economic theory. Our experimental design may include two components: a) an agent-based model (ABM) and b) a laboratory where we carry out a real-world experiment in a typical experimental economics setting. In the ABM we code the rules governing the rational behavior of each individual agent, and by running the model we observe the patterns of global economic behavior arising in the agents' community. In the experimental economics laboratory we explain to the subjects the purpose of the experiment and let them perform economic transactions which we also observe and record.

In the modeling literature, if some regularities are found in the behavior displayed by the agents' community as a whole, it is often said that these patterns have "emerged" from the agents' interaction. If these patterns are similar (according to some problem-dependent criteria) to the ones displayed by the subjects in the real-world experiment, we say that the model has been validated, at least to a certain extent. The approach I have described so far is commonly accepted in the complex systems science literature.

Let's now suppose that, *while* we run the experiment, a colleague develops a new hypothesis, according to which if agents are reminded of ethical values prior to executing an economic transaction, their behavior may change in ways not accounted for by classic economic theory. The research team is intrigued by this suggestion and decides to test it: From a certain time onward, agents are reminded of ethical values although not directly instructed to act differently. This action is equivalent to an intervention in the experimental setting, that is, in the system under study. In accordance with the original purpose of the research, a similar intervention needs to be imposed on the ABM in order to carry out a comparison between modeled and observed results. The question I address in this paper is how such an intervention can be carried out in the ABM.

Attempting to answer this question leads me to suggest that only certain processes can be modeled, and that the widely accepted assumption that in principle most natural processes can be modeled is based on expectations which today we are not able to validate scientifically.

2 Scope of the analysis

The thought experiment I described above implies certain statements and specific concerns which define the purpose and scope of this paper. I make them explicit in this section and I then analyze them in more detail in the remainder of the paper.

Firstly, as a model, I define any formalization of a process, whether via a computer program that generates a simulation, a closed-form mathematical description or a formal logical system. In Section 3, I justify this definition by highlighting the analogies between these different representations.

Secondly, I am interested in the emergence of processes or behaviors with *causal* power, that is, of processes that can *act* (or appear to be able to act) on other processes. In other words, in this work I am not interested in pattern formation, that is in regularities that can be exploited by an observer to deduce certain system properties. I discuss this briefly in Section 5, while I refer to (Boschetti et al. 2008) for a more in-depth discussion and to (Boschetti and Gray 2007b) for a discussion of our ability to discriminate effective versus apparent causation.

Thirdly, I associate causation with intervention; I do so not only to circumvent deeper philosophical challenges, but also because I am interested in learning where, when and how I can intervene in a system in order to understand it, control it or manipulate it.

Finally, as already implied in the approach taken towards causality, the purpose of this work, despite being apparently abstract, is extremely practical: I am an applied scientist who uses models to understand real, complex social and ecological systems. In these systems real-world experiments can be very expensive, time consuming, ethically questionable or simply impossible and thus are increasingly replaced by *virtual* experiments carried out via modelling. Understanding to what extent these experiments are informative in enabling us to understand and predict emergent processes is of crucial importance to problems of immense social relevance like climate change, economic stability and global sustainability.

3 Computation

Computation is relevant to the discussion on causation and emergence at least at three levels. Firstly, there is an equivalence between computation and formal logic (Chaitin 1997). They both start from some fundamental set of axioms and rules and they both generate outputs that are obtained by transforming the a priori set via the rules.

Secondly, formalization is often seen as the ultimate achievement of scientific inquiry, the use of mathematics in physics being the best example.

Finally, many scientists use numerical models to carry out computations in the form of simulations of physical, biological or social processes. For the working scientist this is seen as a pragmatic approach, somehow opposed (and “inferior”) to rigorous mathematical analysis: If we could write and solve closed form equations we would not build models that approximate the same process numerically. However, according to the equivalence mathematics=formal system=computation, the distinction between numerical models and closed form equations is irrelevant in the context of our discussion.

4 Causality

At a practical level, the common challenge causality gives the working scientist lies in its discrimination from correlation: Given two correlated events *a* and *b*, we often need to ask whether *a* causes *b* or whether their correlation is due to a hidden common cause *c*. When possible, an effective way of addressing the problem is via intervention: By imposing a chosen perturbation on event *a* and observing the consequence on event *b* we may be able to unravel the underlying causal structure (Pearl 2003).

In the thought experiment above an example of intervention was given as the act of reminding the subjects of ethical values. The question we asked was how an analogous intervention can be carried out in the model. The question can be rephrased as asking where the experimenter should intervene in order to change the behavior of the agents.

As we saw in Section 3, in a logic system everything can be inferred from the axioms and transformation rules,[i] which in a model are represented by the input and the algorithm, respectively. Consequently the only avenue to changing the model behavior (to intervene in the system) is to modify either the input or the algorithm. Thus, the obvious answer to the above question is that the model should be stopped and re-written by changing or enlarging the agents’ rules; once the model is launched, no intervention is possible to change the agents’ behavior.

We can conclude that causation (understood as ability to intervene) in a formal system should be ascribed only to axioms and fundamental laws; analogously, in a computer model causation can be ascribed only to the input and the algorithm. As a consequence, by construction, all events *a* and *b* in a computer simulation that are not part of the input or the algorithm are merely correlated and can not be in a causal relationship to one another.

A similar discrimination is discussed by (Rosen 2001) under the terms ‘causal’ and ‘logical entailment’; causal entailment implies the effective causal power of agents or physical laws to affect other processes, while logical entailment implies the mechanical inevitability of events related by logical necessity.

5 Emergence

Pragmatically, the working scientist faces emergence when a system displays different behaviors at different levels or scopes (Ryan 2007) and he or she needs to identify how to study them, how they relate or how to act upon them. For example, social systems affect and are affected by processes at the level of an individual’s psychology, economic expectations, family relations, social values, political environment and underlying culture, to name a few. Shalizi formalizes this view via the concept of efficiency of prediction: “[A] feature is emergent if it can provide better predictability on the system behaviour, compared to the lower level entities” (Shalizi 2001).

Since I am not merely interested in prediction, but in intervention, I modify Shalizi’s definition roughly as follows: A feature is emergent if it displays causal power in terms of intervention.[ii] This matches almost exactly the argument proposed by Pattee in addressing the dilemma of where causality lies in Nature: Causality should be attributed to the level(s) at which we can intervene (Pattee 1997). In other words, according to Pattee, while it may be true that all natural processes may ultimately be ascribed to the laws of subatomic physics, for most practical purposes these laws have no effective causal power for us as agents since we can not manipulate them. To intervene in most processes of human relevance we need to find different levels or processes over which we can impose or exploit causal power; these are the entities we can consider as emergent.

The relation I suggest between emergence, causal power and possibility of intervention highlights a crucial difference between a real-world process and its representation as a model. In a model, logical entailment (the only entailment possible in a formalization) places causation solely at the level of fundamental generative elements (axioms and laws), that is, only at the “lowest” level. In a real-world system, identifying causality with control in favor of causal entailment (the only entailment of practical relevance to agents), places causality at levels *other than* the lowest, since we may potentially intervene at several levels but not at the one of fundamental generative elements. Thus, studying a natural process via modeling and via real-world experiments appear to be not analogous approaches but complimentary ones.

6 Implications of formalization

Hume’s original discussion of causation was based on the reasoning that if we can conceive multiple possible effects arising from the same cause then the relation between causes and actual effects needs to be verified experimentally (Owens 1992). The same concern does not apply to a formal system. In a formal (thus closed) system, there is no space for conceiving multiple outcomes: Events either happen or not, and when they happen they do so in a manner rigorously prescribed by the logical entailment originating in system axioms and rules; conceiving a different outcome is equivalent to making a “wrong” logical inference.

The act of modeling a natural process thus corresponds to converting causal entailments into logic ones. This has a number of consequences: Firstly, converting a causal entailment into a logical one effectively eliminates the need for empirical verification of causation since by construction, in the conversion, causation is moved into the axioms and rules. Secondly, when such axioms and rules represent the “lowest” level of representation (as it is most often the case) formalization also prevents downward causation as well as the emergence of casual power at any “higher” level (Boschetti and Gray 2007b).

7 Unreasonable expectations?

The scientific method is based on experimentation, formalization and falsification: The first and third items uniquely distinguish it from other forms of human knowledge. Formalization is not the essence of science, but rather one component, a working framework of pragmatic use, whose deductions always need experimental validation, questioning and falsification.

However, in the discussion on emergence and causation, formalization often plays a more essential, almost dogmatic role: Since causation does not arise explicitly from fundamental physical equations its existence is questioned; since the arrow of time (closely related to causation) is not needed in fundamental physical equations the existence of time is questioned; since emergence of casual power and downward causation are non compatible with formalization, they also are questioned.

Humans have a strong perception of time, causality and emergence; how does this affect the scientific discourse? On the one hand this is often held as a justification for not abandoning these concepts in practical work; on the other some theorists argue that they merely represent an anthropocentric view of nature, unrelated to its real inner workings.

It is reasonable to ask whether there is any *scientific* (that is, falsifiable) reason why the anthropocentric view should hold less significance than formalization in the scientific discourse. I rephrase the question by asking: What are the conditions under which it is reasonable to discard causality as well as emergence, thereby giving priority to formalization over perception? Or in the view of the previous discussion: what are the conditions under which it is reasonable to discard causal entailment for logical entailment, as the only process occurring in Nature?

It seems to me that a possible answer needs to be based on the following four conditions:

1) Nature can be formalized (at least in principle);
2) such formalization is accessible to us;
3) we have shown that such formalization is available for certain natural processes; this leads us to trust that
4) all other natural processes will sooner or later also be formalized.

It seems to me that today the above conditions can not be held as scientific facts. Conditions 1 and 2 are probably beyond verification (Wolpert 2002); condition 3 offers *satisfactory* experimental validation in few fields; condition 4 today is only an expectation.[iii] Scientific knowledge should not be based on an expectation.

In other words, I suggest that replacing causal entailments with logic ones, effectively assuming that all natural processes including the emergence of causal power can be modeled, demands an unreasonable expectation on our current scientific knowledge and formalization. Such an expectation is not falsifiable and consequently it is, strictly speaking, beyond the reach of the scientific method.

8 Implications for research

Hume’s skepticism, intervention as a recipe to identify causation and effective emergence, and the incongruity between logic and causal entailment seem to direct us toward a “back to experiments” approach to complex system science rather than towards more formalization. Still, as discussed above, certain “complex” experiments are difficult or impossible to perform and need to be carried out via the “virtual” analogues of computation; but computation, also as discussed above, prevents the emergence of causal power, which in many complex cases is the focus of our interest. Does this not defeat the purpose?

We need to distinguish two approaches: One involves using a numerical model to study the unexpected consequence of the causal relations we formalize in the code. This seems to me a logically admissible use of a numerical model. The second involves using a model to infer causal relations. This requires intervention and assumes that formal systems can generate causal laws beyond what is embodied in the axioms and generative rules, which is not possible (Boschetti and Gray 2007b).

Things change considerably if the computational experiments allow for the modeler’s intervention in the running of the code, or formally acknowledge the action of rewriting the code as part of the experiment, or both (Boschetti et al. 2008). In the thought experiment described in Section 1 one component includes the ABM, which is a computer model; in fact, the ABM can not exist in isolation from the modeler. In a system that comprises the ABM and the modeler, some form of causal entailment is possible in the action of the modeler who intervenes by altering the code or the model running.

This may happen in two different ways. The modeler may *interact* with the model by providing input *while* the model runs:[iv] In the thought experiment described above, this may allow the modeler to change the way agents carry out economic transactions. However, for this to be possible, the alternative agents’ behavior needs to be already coded, as does the interface allowing the modeler to intervene. In other words, the intervention needs to be envisaged a priori, before the running of the code. For all practical purposes this prevents such a model from being built, since it is unreasonable to expect that the modeler foresee all possible interventions we may decide to carry out in the future. This leaves as the only feasible option the necessity of the modeler to stop the model running and recode the model when a new process needs to be studied.

This second option requires a number of decisions from the modeler. Can the new process be modeled by intervening at the level of the existing rules and input? Or should new rules be included which act at a different level? If so, how will the different rules interact, that is, how will the code routines nest into each other? Most importantly, will the rules at different levels be able to coexist in the code? These questions highlight the potential existence of different levels of code rewriting and adaptation to user interference, which may relate not only to the functionality of the code but also to the nature of the process under analysis. I believe that a careful analysis and a classification of what type of causality can be analyzed via computation as a function of these different levels of code rewriting is of crucial importance for the advance of complexity science and for our understanding of causality and emergence in real-world processes for which traditional experimental work is not possible. This is the avenue for further research that I intend to follow.

References:

Boschetti, F. & Gray, R. (2007a). Emergence and computability. *Emergence: Complexity and Organization, 9,* 120-130

Boschetti, F. & Gray, R. (2007b). A Turing test for Emergence. (In M. Prokopenko (Ed.), *Advances in applied self-organizing systems* (pp. 349-364). London: Springer-Verlag.)

Boschetti, F., McDonald, D. & Gray, R. (2008). Complexity of a modelling exercise: A discussion of the role of computer simulation in complex system science. *Complexity, 13,* 21-28.

Chaitin, G. (1997). *The limits of mathematics: A course on information theory & limits of formal reasoning*. (Springer, New York)

Owens, D. (1992). *Causes and coincidences*. (Cambridge University Press, Cambridge)

Pattee, H. (1997). Causation, control, and the evolution of complexity. (In P. B. Andersen, C. Emmeche, N. O. Finnemann, & P. V. Christiansen (Eds.), *Downward causation* (pp. 322-348). Århus: University of Århus Press.)

Pearl, J. (2003). Statistics and causal inference: A review. *Test Journal, 12(281),*345

Rosen, R. (2001). *Life itself: A comprehensive inquiry into the nature, origin, and fabrication of life*. (Columbia University Press)

Ryan, A. (2007). Emergence is coupled to scope, not level. *Complexity, 13,* 67-77

Shalizi, C. (2001). *Causal architecture, complexity and self-organization in time series and cellular automata.* (Santa Fe: Santa Fe Institute)

Wolpert, D. M. (2002). Computational capabilities of physical systems. *Physical Review E, 65,* 016128.

---

[i] I purposely avoid any incompleteness argument in this draft; for an analysis of incompleteness in causality and emergence I refer to (Boschetti and Gray, 2007a) and references therein.

[ii] I do not intend to suggest this as a formal definition of emergence, since it is neither novel nor general, but rather to use it as a pointer in this discussion.

[iii] Note that the belief that all natural processes can be reduced to a handful of sub-atomic physical laws is based upon this expectation.

[iv] This is not accounted for in classical computational theory, since the classical Universal Turing Machine employed in this field of research does not allow for external intervention during its running. This distinction has potentially important theoretical consequences that go beyond the scope of this paper.