HPLS (2014) 36(1):60–89
DOI 10.1007/s40656-014-0005-2

ORIGINAL PAPER

Mapping an expanding territory: computer simulations in evolutionary biology

Philippe Huneman

Received: 28 November 2013 / Accepted: 16 March 2014 / Published online: 25 July 2014

Abstract The pervasive use of computer simulations in the sciences brings novel epistemological issues discussed in the philosophy of science literature since about a decade. Evolutionary biology strongly relies on such simulations, and in relation to it there exists a research program (Artificial Life) that mainly studies simulations themselves. This paper addresses the specificity of computer simulations in evolutionary biology, in the context (described in Sect. 1) of a set of questions about their scope as explanations, the nature of validation processes and the relation between simulations and true experiments or mathematical models. After making distinctions, especially between a weak use where simulations test hypotheses about the world, and a strong use where they allow one to explore sets of evolutionary dynamics not necessarily extant in our world, I argue in Sect. 2 that (weak) simulations are likely to represent in virtue of the fact that they instantiate specific features of causal processes that may be isomorphic to features of some causal processes in the world, though the latter are always intertwined with a myriad of different processes and hence unlikely to be directly manipulated and studied. I therefore argue that these simulations are merely able to provide candidate explanations for real patterns. Section 3 ends up by placing strong and weak simulations in Levins' triangle, that conceives of simulations as devices trying to fulfil one or two among three incompatible epistemic values (precision, realism, genericity).

Keywords A-Life · Computer simulations · Explanations · Models · Evolutionary biology · Hypothesis testing

In the spirit of logical positivism, it would not be unusual to think that the two sources or "ingredients" of scientific knowledge are empirical evidence gained

---

P. Huneman (✉)
Institut d'Histoire et de Philosophie des Sciences et des Techniques, CNRS/Université Paris I Sorbonne, 13 rue du Four, 75006 Paris, France
e-mail: philippe.huneman@gmail.com

through experiment or observations, and logic—especially in the form of mathematical models. The eruption of computer modelling changed the picture of science to the point that what now may be the most important part of the daily work of scientists—especially in the fields of the special sciences—consists in designing and running computer simulations; it took a few decades before philosophers of science began to take such a change seriously and to wonder whether scientists were still doing the same science as the one the Vienna Circle was philosophizing about. In the wake of Humphreys (2004), they thus began to ask how we know something through computer simulations, and, more radically, what do we know through them (Winsberg 2010). If one can agree that simulations extend our computational abilities in a way somehow parallel to the way instruments such as telescopes or microscopes extend our perceptual abilities then it is understandable that philosophers should assess the extent of the consequences of such an extension. Computers do not only imply that we compute faster and more: according to Anderson's famous slogan ("more is different", Anderson 1972), they also imply that we compute, and then think, differently. The philosophical question thus becomes: how different? Are computers a new sort of experiment? (Peck 2004).

In the context of these questions it is interesting to consider one specific kind of computer simulations, namely those that are used in evolutionary biology, for the following reasons. Besides the fact that a huge amount of models and theories are induced and tested on the basis of computer simulations, it is notable that a whole field of biology deals only with computer simulations: so-called Artificial Life (AL), which tries to design algorithms that implement the properties of inheritance, variation and fitness likely to yield a real evolutionary process. The philosophical rationale behind this project is often the desire to investigate the proper features and processes of life—assuming that life is what evolves, on Earth, through evolution by natural selection—as they would be when separated from the fact that living individuals on Earth have been instantiated on the basis of a carbon-based chemistry. This is arguably a contingency of our universe; and if "Life" is supposed to make sense in general—that is, to be something like a "natural kind" (sensu Kripke) or to be able to refer across possible worlds as a rigid designator—one should be able to disentangle the properties of Life itself (or at least of the agents likely to undergo evolution by selection) and the properties stemming from the fact that carbon molecules instantiate life in this world.

Moreover, turning to evolution itself, it has even been argued (e.g. Dennett 1995) that evolution by natural selection is a sort of algorithm. Inversely, in all areas of science and engineering, when it is not possible to solve an optimizing problem as a solution of equations, one designs a specific range of algorithms, so-called "genetic algorithms". Their description uses the keywords of evolutionary biology—genes, mutations, recombination, crossover, fitness, etc.—which also indicates that evolution may receive an algorithmic interpretation. When one adopts such interpretation, it is natural to think that the computer simulations invented and explored by AL scientists capture something of the algorithmic processes that are constitutive of evolution. It is therefore plausible to think that the epistemic issues proper to the use of computer simulations can be fruitfully addressed within the context of evolutionary biology. Examining how these simulations are designed, the questions they are likely to answer, and the kind of explanations or contributions to explanations they provide should enrich a general understanding of knowledge by simulations. This is the topic of the present paper.

The outline of this paper is the following. A first section will expose the specific epistemological problems raised by computer simulations, namely, the validation problem, the explanatory scope problem, and the problem of their relations to non-simulation models. Essentially, simulations will be examined alongside the antagonistic epistemic values that Richard Levins elaborated with respect to model building around 50 years ago. The second section will sketch some roles of simulations in evolutionary theory. The last section will isolate what is epistemically proper to the use of simulations in evolutionary biology, and the types of responses to the epistemological problems outlined in Sect. 1.

1 Computer simulations and their epistemological problems

1.1 Characterization of computer simulations

Provisionally, I work with the following characterization of a computer simulation in general: simulations include programs which drive the behaviour of a set of digital elements with specific rules—such as agent-based models, cellular automata (e.g. Margolus and Toffoli 1987), and genetic algorithms, all of them being used in all areas of the sciences, natural as well as social (see, e.g., Epstein 1999). In doing so, they may represent a target system. For instance, an agent-based model simulating the dynamics of a deliberating group—whose rules (either deterministic or stochastic) are founded on the changing beliefs rules inspired by the psychology of daily life, and whose prior conditional beliefs are based on average opinion polls gathered from typical representatives of the given sociological categories to which the individuals belong—will model the target system made up of a set of jurors judging a given case defined by some judiciary or social criteria.[1]

But simulations also include algorithms that compute the solutions of differential equations once they are discretized. For example, while general Navier–Stokes equations capture the behaviour of liquids, they are hard to solve generally; and therefore their discretized counterparts allow one to compute approximate values of the solutions. These two types of simulations can thus be embraced in a single concept: a simulation realizes an algorithm implementing a specific model (as a simplified representation of a possible state of things), which can be defined by discrete rules or by equations, into a computer language in a way that allows us to compute a set of outputs from a set of inputs.

For a given simulation, then, running the program—which can be called a "simulation round"—takes the place of deducing the consequences from an analytic model, i.e. computing the integrals or solving a set of partial differential equations. The algorithms are determined by parameters whose values range across a particular domain. Comparing identical and different rounds of the simulation (with regard to

---

[1] See http://www.jurysimulationresearch.org to find examples of such simulations.

parameter values) allows one to "gauge" the space of parameters according to the kinds of outcomes each combination of parameter values yields. In the sense that it allows one to capture the possible behaviours of the target system, this operation precisely corresponds to solving the equations of an analytic model.

1.2 Three questions

Once simulations are understood as models, they raise epistemological questions along the following three lines: their explanatory scope, their relation to experiment and mathematics, and their validation.

1.2.1 Explanatory scope

Even though "model" is a term at the centre of heated controversies in the philosophy of science, almost everyone agrees that models play an important explanatory role in general. Computer simulations are a specific kind of model, but there is not a universal answer to the question: "what do they model?" For example, consider the famous simulation of the flocking behaviour of birds proposed by Reynolds (1987). This simulation showed that digital creatures called "boids"—obeying very simple local behavioural rules of direction matching, collision avoidance, flock cohesion, may display a pattern of collective motion similar to the patterns of collective flight displayed by flocks of birds.[2] "Boids" are also a good model for the behaviour of schools of fish and other groups of animals.

However, nothing in the model transcribes the real characteristics of fish or birds: the only property under focus is that they move. No knowledge of what fish or birds are is relevant here—not even the fact that they are living organisms, or that they evolved through natural selection. This is in contrast with other kinds of models. Experimental models—think of Mendel's peas—model systems of populations with inheritance properties, since they are themselves sexual species. It can hence be seen as a legitimate model of biological inheritance. Mathematical models, on the other hand—think of Fisher–Wright models in population genetics—are modelling populations of genes, some main dispositions of which (mutating, recombining, replicating, impacting reproductive chances of phenotypes) are cardinal features of the entities called "alleles" in the model. It can hence be seen as a legitimate model of allele frequency change in a population.

In the absence of references to biological entities or processes, it is not obvious that the "boids" simulation, for example, is a model of animal behaviour: boids do not share (like peas) any biological property with any animal system; and boids are not modelling some of the main dispositions common to all the "animals" *qua* animals they would model—such as fish or birds. Granted, the explanatory relation of models to reality—or of "model systems" to "target systems" as this is called in the literature—is a longstanding puzzle; however, nothing indicates that such a puzzle will be the same when one considers computer simulations; therefore, it may

---

[2] The simulated flocking behaviour can be seen at http://www.red3d.com/cwr/boids/applet/.

be that the boids model does not exactly pertain to biology the same way as the other two examples—Mendel's peas and the Fisher–Wright equations—do.

1.2.2 Relation to experiments and mathematics

The boids simulation is both a digital system, and an experimental setting—like Galileo's inclined planes, with which it has many features in common, as it has been often argued in the literature (e.g. Peschard 2010; Hughes 1999; Norton and Suppe 2001). It is a simple system that produces an output according to how we determine the input: stones of various weights and various inclination angles, in the case of Galileo's device; various parameter values for the rules of the boids, the number and different initial speeds of the boids, the positions of obstacles along the trajectory. In both cases, we gain knowledge by considering the various reproducible associations between our chosen inputs and the recurrent outputs. Moreover, this knowledge is about something else than the model system that is somehow related to it: the dynamics of gravitational bodies on Earth, the swarming behaviour of animals. Along these lines, it has been suggested that computer simulations are indeed a sort of experiment, using digital material rather than matter (Morrison 2009).

It is quite common to think that in many cases simulations are designed when experiment or observation is not possible (claim 1); climate science or cosmology provide instances of this. However, the fact that computer simulations simply implement some algorithms also inevitably suggests that they are mathematical tools rather than experiments. Therefore, nothing really specific about simulations would distinguish them from mathematical models. Supporting this claim, Frigg and Reiss (2009) argued that there are two epistemic questions which scientists handle: when they use simulations as well they work with mathematical models checking the appropriateness of the model ("validation") and verifying its computation of the solutions ("verification").

It is indeed common among scientists to view simulations, which are algorithms, as a mathematical alternative to models written with equations. As a result, the question for them turns to the choice between these two kinds of mathematical models: authors debate about which criteria should govern this choice; and further, whether there are ideal recommendations about which should be chosen in general. And here, a plausible view is that simulations are used when the solutions to the analytic models are not available (claim 2). Claim 1 and claim 2 therefore see simulations as a *surrogate*—whether the inaccessibility concerns experiments or equations. It is possible that both are right regarding some kinds of simulations.

However, with regard to the relationship between analytic models and simulations, three cases (i–iii) have to be distinguished:

(i) We may have an *analytic model* of a system written in terms of equations that describe its behaviour. Recall the Navier–Stokes equations which govern any hydrodynamic system, and would therefore also govern for example its turbulences. It is in general too difficult for us to find the solutions of these partial differential equations (which would allow us to describe a given hydrodynamic system, model its behaviour, and predict phenomena of interest). Therefore, we simulate the system under focus based on some algorithms derived from the Navier–Stokes equations. Often, we may have no actual data to compare to the outcomes of the simulation, since simulations are used when we have no access to the reality modelled by our analytic models (think of simulations of the behaviours of galaxies when they crash into each other, black holes, etc.).

(ii) Let us suppose there is a system for which we have *no equations.* Yet we may have some coarse-grained description of it—some reasonable intuitions about things that are going on that allow us to build a model by considering the behaviour of parts of the system along some rules we intuit, and then modelling the global behaviour and its outputs. Agent-based modelling often proceeds in this way; it is thus always bottom-up modelling. An example would be Reynolds' boids—since scientists initially thought of some possible rules that bodies moving in a flock- or swarm-like manner could follow, built the boids accordingly, and then fine-tuned the rules and checked for the accurate output (the one that corresponded to empirical data). Many models in the social sciences are like this, for example, Epstein's model of civil violence (2002). These models correspond to what Humphreys called "non-theory based models". In this case, we are not facing a situation where simulations are used because analytic models are intractable, since these models are not available from the start.

(iii) There are also cases where the same phenomenon can be understood *both through an analytic model* (i.e. partial differential equations) *and through algorithms*. More precisely, it is often the case for certain kinds of systems that we have an analytic model where some subclass of behaviours of the system is tractable while many others are not—while we have a range of simulations for any subclass of behaviour. Simulations and equations may be very different, with no easy way to derive the latter from the former; and they may also have different simplifying assumptions, which implies that they are not exactly equivalent. Sometimes this case includes situations where researchers start with the simulation models. These simulations generated interesting outputs that proved they were a correct model for the data. Afterwards, some researchers found the analytic model corresponding to the simulations. This happened with the neutral theory of ecology—first exposed in the seminal book by Hubbell (2001) in terms of simulations of agent-based zero sum games in metacommunities,[3] with the agents being individuals of specific species. Later on Volkov et al. (2003) derived a set of equations that generally described these neutral metacommunities. However, these equations are not easily tractable.

Case (iii) raises the question of how to choose between equations and algorithms—a methodological issue often faced by scientists, as is exposed by Gaucherel et al. (2011) in the case of ecology. They argue that there is no principled criterion in

---

[3] Sets of ecological communities of species, related by processes of migration, colonization, gene flow, etc.

choosing one above the other; rather, complex explanatory interests dictate preferences. Nevertheless, the alternative between analytic models and simulations is often understood in the sense of a lesser value of simulations—because either they just approximate the solutions, or they are simply embedded in gross common sense assertions or intuitions—a fact that lowers their scientific value.

The contrasting reactions to computer simulations in explanations range from depreciation by microeconomists (Lethinen and Kurikovski 2007) to enthusiasm among the community of Artificial Life researchers (Adami 2002). This variety has to do with a paradox that emerges when we compare modelling by simulations with experiments and mathematical models. If we consider simulations as experiments, they have an original characteristic that could be called *transparency*: we are somehow in command of the simulation, since we control all its parameters; whereas with experiments, which take place in the real world, we may fail to know important parameters or causally relevant factors that are accountable for the outcomes.[4] If we consider simulations as mathematical models, on the other hand, they are non-transparent, in the sense that we often cannot write down a solvable equation that would account for their global behaviour; in other words, these outcomes are not necessarily solutions of analytical equations of some canonical type. For instance, in many cases of cellular automata we cannot look globally at a step of a cellular automaton and compute its next step without looking at it locally, cell by cell (Wolfram 1984; Huneman 2008a).[5]

The reservations vis-à-vis simulations may often support the quest for an analytic model that would transcribe the algorithms. Turning again to Reynolds boids, one notices that this is actually the case here. The boids algorithms have been very illuminating in regard to the behaviour of gregarious animals, such as nightingales or cranes; it also perfectly models the behaviour of schools of fish.[6] It shows that simple individual behavioural rules, which only consider the local position and neighbourhood of a fish or a bird, generate a cohesive and ordered motion without the need of any centre of control. Later on, it was shown to be derived from the agent-based model of self-propelled particles (SPP) called the Vicsek model (Vicsek et al. 1995)—where particles follow discrete rules that do not consider long-range parameters (as was the case in Reynolds rules: they fly in accordance with the centre of the group…), and noise is added (which plays a fundamental role in making the collective behaviour robust; see Toner and Tu 1998). SPP are a paradigmatic model for the study of swarming behaviour in all unicellular organisms (to which no sensation of long range events can be ascribed) and animals,

---

[4] Often, unpredicted results of experiments, or failure to reproduce experiments in different labs, come from these parasitic unknown factors: most recently, the European experiments that found neutrinos going faster than the speed of light were affected by such uncontrolled factors, i.e. optical fibers with unusual diameter (Reich 2012).

[5] Moreover, computer simulations run by themselves, so we cannot see and understand what is going on—in the sense of how Descartes, in principle, required all cognitive operations to be self-certain of the validity of each of its steps.

[6] It is still used when producing cartoons with fish [Reynolds' simulations were first developed in the context of video programming, as is indicated by the fact that his paper was published in a journal of graphic design (*Computer Graphics*) rather than in a computer science or biology journal].

including humans. And in turn, it has been shown that the Vicsek model instantiates a set of equations known as the Toner–Tu equations—which describe the whole flock behaviour through a set of four equations with two variables, the flock density and the flock (vectorial) mean velocity (Toner and Tu 1998). And so here we have a continuous epistemic process that went from simulations like case (ii) to simulations like case (i) (equation-based).

Therefore, it looks like once again, as in the case of the neutral theory in ecology, one can understand the simulation algorithms as the manifestation of underlying equations—just as the zero-sum games meta communities can be seen as instantiations of the equations proposed by Volkov et al. (2003), or as the agent-based behaviour of some aggregated consumption behaviour can be understood as the solution of a set of equations describing the agents' utility functions (as it is assumed by orthodox microeconomists—Lethinen and Kurikovski 2007).

But in this case, moreover, it appears that the Toner–Tu equations are somehow correlated with the Navier–Stokes equations that describe the motion of a fluid: in both cases, the overall behaviour of macroscopic entities is described through formally analogous equations of velocity and density. Yet, the former equations violate a symmetry principle proper to hydrodynamic liquids in equilibrium described by Navier–Stokes—namely, the fact that such liquids do not have a privileged direction in three-dimensional space, unlike SPP collectives.[7] But the epistemically relevant aspect of this analogy is that these Toner–Tu equations provide a new insight if one considers a striking feature of the collective behaviour of birds viewed in nature: seen from far away they look like the flow of a river. They are especially analogous to this when they avoid obstacles by separating into groups and rejoining afterwards—just like a river moving around a rock or a sand island (see Fig. 1).

Now, imagining that the fish in a school of fish become smaller and smaller as the school becomes larger and larger, the school starts to resemble a flow of particles (like, for instance, iron dust); and at the limit, when the particles reach the size of molecules, the school would become a liquid (ignoring, however, the symmetry breaking proper to flocks; cf. Lopez et al. 2012). Therefore, the fact that the behaviour of boids as modelled by Reynolds resembles the behaviour of a flow of a liquid as seen from very far away is captured by the fact that the Toner and Tu equations governing these simulation systems are a sort of Navier–Stokes equation. Actually, it is "the existence and similarities of collective motion phenomenon in very different species and even in non-living systems" (Lopez et al. 2012) that stimulated the elaboration of a general set of equations, which captures flocking behaviours and more generally SPP behaviour as an instance of a universal behaviour. So here, the simulations can appear afterwards as being derived from equations; but in fact, the equations just provide a supplementary understanding of what is going on in collective behaviours—i.e. the fact that they instantiate a very general behaviour in nature, affecting both organisms, small particles and liquids.

---

[7] Even though the equations do not specify the cause of the symmetry-breaking, and that it can as well be initial stochastic fluctuations.

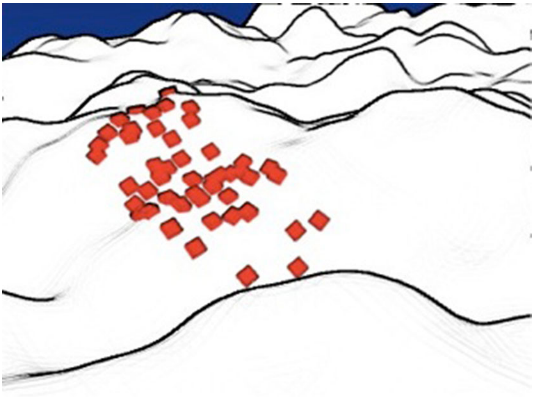

Fig. 1 Reynolds' boids. Screen capture from http://commons.wikimedia.org/wiki/File:Blender3D_BoidsOnPlane.ogg?uselang=fr, which shows the flight of boids over mountains (*source* Wikipedia Commons)

1.2.3 Validation

Simulations may raise specific problems about what they explain, about their intrinsic nature, and—given that they are scientific tools—about their validation. Once again, Reynolds' boids are a good example. As said, they adequately model the behaviour of flocks of birds as well as schools of fish by ascribing to agents in the simulation (the boids) some simple behavioural rules. The result matches the observed patterns of behaviour of many social animal species. However, what do we mean by claiming that this is a good model? Fish, birds, grasshoppers are very different species, and nothing indicates that they may have the same nervous system; therefore, how could you say that they genuinely obey the same behavioural rules? Without independent evidence of the fact that the nervous or cognitive systems of specific fish and birds indeed implement Reynolds' rules, we cannot say that the model accurately captures the nervous and cognitive properties of those animals—even though we could say that we correctly model the behaviour of the collective. It is a mathematical fact that, if these rules can produce the observed flocking behaviour, there are possibly many other sets of rules that yield the same grouping behaviour. Thus, why would we say that all species that display the same behaviour therefore obey the same rules, namely those proposed by Reynolds? Adding to this that the boids are in fact instantiating a Toner–Tu equation that also holds for purely physical systems increases the doubts we could have about the modelling adequacy of the boids. Even though the boids prove that no centred control system is required to produce flocking behaviour in any species, it is hard to infer more than that about the genuine rules that govern individual behaviours.

This validation problem for simulations seems to be quite original, if compared to the usual model validation problems, since the rules of the building blocks—e.g. the agents in an agent-based model—are often not intended to be the rules actually followed by extant parts of the system (i.e. organisms) that could be studied empirically. There is also, however, an analogy with the underdetermination of theory by its data; i.e., as Russell or Duhem remarked a long time ago, any mathematical "law" which is hypothesized to explain observed regularities comes with a whole series of different laws that are empirically equivalent.

To make the validation problem for simulations more salient, as well as to capture its originality, let us turn to meteorology. Küppers and Lenhard (2005) have shown that simulations of the global dynamics of a system sometimes make very unrealistic assumptions in regard to the parameters of the system they model (in order to compensate for the effects of the discretization of the equations), but yield very adequate predictions of global behaviour. Arakawa's climate models are good examples of this: they quite accurately predict the future weather on a global scale, but they set some parameter at a highly unrealistic value. Here, not only are the unrealistic assumptions tolerated in the simulations, but without them the model would not be accurate enough in regard to what we want to predict.

1.3 A general perspective on epistemological questions relating to simulations

A more fruitful way of considering the validation problem, as well as the explanatory scope problem, is to recall Levins' (1966) conception of the strategy of model building. In this 1966 paper, when considering ecological models, Levins argued that models have *different* epistemic goals which cannot be fulfilled at the same time (all of them being sometimes conflated under the name "validation"). Generality, realism, and prediction are these goals: a model can be realist and take into account the specificities of a situation (for example a predator–prey model of a population of foxes and rabbits could realistically describe the dynamics of a population of foxes and rabbits). Although a pure predator and prey model (e.g. Lotka–Volterra equations) will be more general, it loses realism—even though it can give rise to predictions when some of the parameters are informed. Further, by adding more specificities of an extant population, the model can be very precise as it entails reliable quantitative predictions; however, it will lose some generality since other prey and predator populations will not be modelled by it (Fig. 2).

As a consequence, Levins argues that no model can satisfy all epistemic values at the same time, and therefore one has to make choices: there exist trade-offs between all values, and the realization of these trade-offs characterizes a given model. That is also why the same phenomenon, or target system, can be modelled by several very different models that may have very different assumptions. Further research (e.g. Matthewson and Weisberg 2009) delves into the details of these trade-offs by distinguishing and defining different types of models by the kinds of trade-offs they can make.

In the case of Arakawa's climate model, it is easy to understand what is going on: it trades off realism for precision while keeping the level of generality unchanged. The boids are even more general, and possibly less precise. They also forget realistic details. So generally speaking, these models tend to exist in the part of the triangle that is mostly general and precise (defining "precision" by the accuracy of global predictions). Therefore, their validation cannot be understood as a comparison with the real systems they model—but rather as global predictive accuracy.

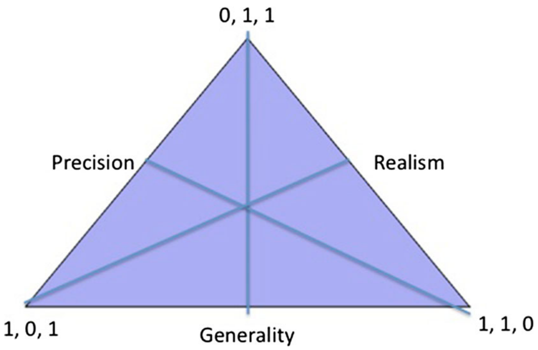

Fig. 2 Levins' triangle. The coordinates are triplets (G, R, P), G being the value for generality, R for realism, P for precision. Each edge defines value 1 for G, R or P (drawing by the author)

2 Simulations in evolutionary biology

There is an extensive use of simulations in evolutionary biology for a specific reason; and this reason allows for some interesting features in these simulations that may provide determinate answers to the three epistemological problems outlined above—especially with my intent of putting them within the perspective of Levins' triangle. First, this reason is the fact that the timescales of evolutionary biology are such that it is difficult to experiment on real-life systems. Mathematical models, on the other hand, are often difficult to handle. For instance, genomes contain thousands of loci; but population genetics can maximally do two-loci models, which is little even if it is reasonable to assume that the evolution of one gene can be understood by a two loci model (Gillespie 2004). Therefore, besides classical experiments with bacteria (Lenski and Travisano 1994), *Drosophila* (Teissier and Lheritier) or mice (Lewontin and Dunn 1960), evolutionary biologists often turn to simulations to capture something of the evolutionary dynamics that are likely to occur.

There is another reason (alluded to in the introduction) to focus on evolutionary theory, namely the Artificial Life (A-Life) research program. Evolution on Earth occurs with the descendants of the first living cells—which results in the contingent properties of a shared genetic code and a carbon-based chemistry. To consider what evolution would look like "as such", as the A-Life project asks, simulations are the appropriate tool, be they sophisticated genetic algorithms like Holland's Echo, cellular automata like Conway's "Game of Life", or agent-based models, or any other kind of algorithms. I have elsewhere discussed the project, along with the strong A-Life thesis according to which evolutionary biology is a branch of some general algorithmic theory (Huneman 2012a). In order to address the epistemological problems sketched out in the previous section, I will recall the typology of computer simulations in evolutionary biology before sketching the epistemic properties they have in common.

2.1 Types of simulations in evolutionary biology

In evolutionary biology, many issues obviously concern evolution by natural selection. Therefore, simulations can first be classified according to whether they confer a role to natural selection or not: Some simulations define a context where a form of natural selection can take place—the building blocks of the simulation feature variations, a form of inheritance, and can be ascribed a fitness function.[8] Numerous examples exist: from the simulation of the respective impacts of adaptive potential and geographical barriers on the process of speciation (Maley 1998), to the simulation of industrial melanism, by cellular automata that are in competition for the best match to a coloured environment (Maron 2004). But other simulations in evolutionary biology can do without natural selection: for example, the well-known paleontological model of the diversification of clades according to a random walk, which was developed in order to check to what extent patterns of evolution of biodiversity have to be ascribed to natural selection (Gould et al. 1973, 1977; cf. Huss 2008). Some other examples include the models of ecological change in biodiversity in metacommunities without natural selection, as developed by the neutralist theory of ecology (Hubbell 2001).

The epistemic status of simulations in evolutionary biology gives rise to another distinction:

(1) *Weak simulations* Some simulations may be developed to test a theory (and as in case (i) above in Sect. 1, they may be derived from an analytic model) or a hypothesis (possibly expressed in verbal terms). For example, Maron's melanism simulations mentioned above aim at testing the famous selection hypothesis about the changing colour of peppered moths after the industrial revolution (Kettlewell 1955). Yet "testing" may mean two different things:

    (a) one wants to test a hypothesis that can be modelled in an analytic model but is intractable, or that cannot be modelled by equations at all;
    (b) one wants to test a hypothesis about a process in nature for which real data is difficult or impossible to gather, such as what is going on in black holes. Maley's model, for example, tests the hypothesis that geographic barriers are an important driver of speciation; in doing so, the simulation provides the data by which the hypothesis will be judged.

Thus, within the testing process (a process where scientists conceive of hypotheses, design models to implement hypotheses, and derive predictions and compare predictions to data), the simulation sometimes takes the place of the prediction (a), and sometimes has the role of providing data (b). To sum up, in case a, data and simulations are used to test hypotheses. In case b, there are hypotheses, but no data; this is provided by simulations, which replace experiments and observations.

---

[8] These three properties are, as formulated by Lewontin (1970), the conditions for natural selection.

(2) *Strong simulations* Some other simulations are not intended to test any theory or hypothesis; rather, they just explore the possible outcomes of a simple model. For instance, Holland's Echo algorithm models the ecology and evolution of "organisms" that have genes, have chromosomes, and compete. This is not supposed to model actual target systems; and thus, the correspondence with real organisms—or simply the choice of organismal properties—is not a characteristic of the simulation. Scientists explore Echo and its behaviour by varying its parameter values and checking what happens, focusing on the system itself, and not referring to a target system. Qualitatively, Echo displays some interesting recurrent behaviours (such as the appearance of predators, the emergence of mutualisms and parasitism, etc.) that resemble very general traits of real ecosystems; it shows typical ecological behaviours, such as the competition exclusion principle (Holland 1995; Hraber and Jones 1994; cf. Pocheville 2014 on the competition exclusion principle). But it is not really a model of an ecosystem: its unrealistic assumptions concerning reproduction—namely, the absence of reproductive isolation for species—makes Echo a poor model of evolutionary biology (Crubelier et al. 1997). It is not intended to test hypotheses about any extant ecosystem or kinds of ecosystems and their key features (genetic features such as gene flow; ecological properties). More generally, all simulations developed within the framework of Artificial Life are of this nature. (I will consider below to what extent such A-Life simulations contribute to classical evolutionary biology.) Thus, in contrast to weak simulations (1a and 1b), in strong simulations there are no hypotheses and no data to test.

Often, weak simulations of type a will have a predetermined fitness function. Inversely, in strong simulations, no fitness function is defined—which allows the system to evolve in a sense that may resemble open-ended evolution in which the rate, range and nature of adaptations is not bounded. Actually, some strong simulations are designed precisely to see under which conditions open-ended evolution as we see it on Earth could emerge (Bedau et al. 1998; Taylor 2012).

Interestingly, all these simulations feature the "transparency paradox" mentioned in Sect. 1.2.2. First, they are *transparent*, in the sense that nothing else than the designed algorithm that we conceived of is responsible for the properties we are witnessing. Second, they display emergent processes, which may result in some global outcomes that are not potentially included in the initial states, and not even potentially describable in the language intended to describe the initial states—i.e. the agents and their properties (Rasmussen and Barrett 1995). For example, the moving patterns named "gliders" in the peculiar cellular automaton called Conway's Game of Life (Fig. 3), as well as the so-called "glider guns", are such emergent patterns. Emergence here has to be understood as computational incompressibility (Huneman 2008a, b; Bedau 2003, 2008); for this reason, the emergent outcome is indeed *not transparent*, in the sense that no mathematical global laws are likely to be easily formulated to grasp the reasons for its apparition and predict it.

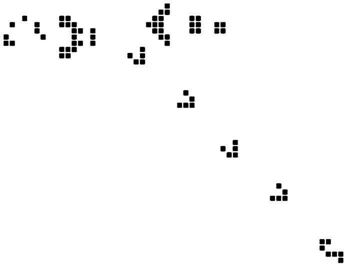

Fig. 3 Gliders in a Game of Life simulation. In this grid, cells are either *white* or *black*, and the state of a cell in line $n + 1$ is determined by the state of the parent cells (*white/black*) in line n and its two neighbors according to a rule. Gliders are these patterns of *black dots* extended through several lines that are conserved as such along many steps of the simulation, therefore that seem to "move" (translate while rotating) regularly through the grid towards the *bottom right*, even though the cellular automaton only determines the state of cells in *each line*

2.2 Epistemic properties of simulations in evolutionary biology

2.2.1 Internal causal dependencies

A general consideration of the epistemic impact of simulations in evolutionary biology starts from acknowledging that they realize a certain kind of causal relationships; and it is in virtue of these relationships that simulations can either represent natural processes or provide knowledge about them, as I will now argue in three steps.

First, let's consider the case of simulations where patterns emerge, such as the gliders mentioned above in the Game of Life: the relationship between the successive positions of the gliders can be understood as a relation of counterfactual dependence between properties of emergent patterns (e.g. "if the glider at step n had not been there, the glider at step $n + 5$ would not be there, and be where it is"). In this sense, Huneman (2012a) argued that in computer simulations in general one can discover counterfactual relations between emergent entities defined at a higher level than the building blocks (such as the gliders in a Game of Life).

Second, and more generally now, the whole class of cellular automata displays counterfactual relations between properties at various steps of the simulation that are definable in terms of disjunctions of possible sets of states of the Cellular Automaton (Huneman 2008a). We assume here that the same kind of reasoning will hold about agent-based models and other types of simulations.

Third, if many computer simulations generally display these kinds of internal causal relationships (counterfactual dependences between global states of the simulation), what about the four classes of simulation in evolution defined above (weak simulations with or without selection, and strong simulations with or without selection)? Considering the case of genetic algorithms (which belongs to the class of "formal selection" contexts simulations), it is easy to see that these causal processes are in fact actual "selective processes". For other simulations in the "formal selection" class, having shown that they undergo some causal processes allows one to think that these processes will indeed be in the form of selective processes. As to the “no selection” simulations, things seem less straightforward and we can leave aside for now their status regarding causal relationships.

Considering the second epistemic distinction (weak/strong), I would now propose the following preliminary distinction about the ways they respectively realize causal relationships:

Weak simulations: If the hypothesis embodied by the simulation is correct, these causal relations are of the same kind as some relations in the target system, e.g. there is an isomorphism between selection in the Maron model of industrial melanism and selection in a real melanism case.

Strong simulations: some features of causal processes occurring in them are isomorphic to generic features of some kinds of causal processes within a class of systems. For example, parasitism as a causal process in Echo does not correspond to some particular parasitism case in a given ecosystem; rather, it instantiates some very general features of the kind of process called “parasitism” (e.g. it may satisfy the models that describe parasitism in terms of cost-benefit in fitness, see e.g. Archetti et al. 2011).

2.2.2 “Pure possible processes” as the representative core of simulations

Having acknowledged causal features of computer simulations, I will now draw some consequences about the epistemic relevance of weak and strong simulations. I start by giving some examples and then defining an original character of evolutionary biology simulations as including causal relations, and then hypothesize what in principle we can expect from evolutionary simulations, especially as compared to *bona fide* causal explanations. On the basis of this insight, the last section will develop answers to the philosophical questions regarding computer simulations in the context of evolutionary biology.

To illustrate what will appear as the main epistemic property of these simulations, let us first consider that you do not have species easily in Echo (Crubelier et al. 1997) or that no lineages can be defined in Tierra. More generally, if one system is designed to study some level of biological reality, the other levels are not *ipso facto* given—whereas they are in real biological systems (e.g., with organisms, you have genes and species). Moreover, as Eldredge (1985) indicated, all biological entities are distributed in two hierarchies: one is defined by *ecological* interactions, since the levels are levels of ecological interactions (cells communicating with cells, organisms chasing organisms, demes or colonies competing for resources); the other is defined in *genealogical* terms—such as genes, organisms, and species—with each level defining a kind of transgenerational persistence in time. In both hierarchies, each level in the hierarchy coexists with all the others.

As a result, in biology it is hard to consider causal relations and interactions at a single level: any cell interaction involves organisms, etc. This is why Huneman (2012b) argued that any computer simulation displays “pure possible processes” concerning the modelled entities, located at a target-level of the biological hierarchy (there is no implicit entangling between levels). By “pure possible processes”, one means that (i) these are “processes”, since there are causal processes at stake in the simulations (as shown earlier); (ii) "pure" because they involve only one kind of entity since the others are not defined in the simulation; (iii) "possible" because one does not know whether these processes actually take place somewhere, or whether they are just potentialities proper to the type of entity modelled in the simulation—which of course relates to the validation problem proper to simulations, described above.[9]

This feature relates to the epistemic property above called *transparency*: the modeller controls all the causal processes occurring. For instance, in an Echo simulation dealing with ecological relations between organisms, no parasitic effect of species–species interaction (e.g. species selection) intertwines with the modelled processes, since species are not even defined. Whereas in real biology it may be that the interactions we witness and conceive of as organismic interactions (e.g. selection of some type of traits) are yielded by a process taking place at another level (e.g. species selection).

In the case of evolution by natural selection, there may be selection at all levels of the hierarchy: genes, organisms, groups, species, and clades (Okasha 2006; Damuth and Heisler 1988). Thus, in the real world, all these selective processes are taking place together (yet with various relative impact)—whereas simulations may model the pure possible selective processes at a given level, e.g., pure species selection in some cases of studies on macroevolution (Gould and Lloyd 1999). If a simulation designs organismic or genic selection but no species is defined, then it shows what selection can do at levels under the species level. Yet of course the "pure possible processes" should not be understood only as *particular-level* selection processes: for instance, the Miller and Todd (1995) simulation considers the pure process of sexual selection and explores its various outcomes in terms of diversification.

Therefore, one can generally see what weak simulations in evolutionary biology, at least, can actually do from an explanatory standpoint: in the "formal selection" context, they test hypotheses regarding the sufficiency of organismic (or species, or genic, etc.) selection to produce a given outcome (e.g. a specific adaptation such as a mating behaviour, or different patterns of diversity). They may also test hypotheses concerning the impact of certain parameters in driving one pure selection process towards a given pattern or outcome. Maley's model, for instance, compares the relative influence of adaptation potential and geographical barriers on speciation as a possible outcome of a pure organism-selection process.

Even if we turn to "no selection" simulations, an analogous feature appears. For example, the Gould et al. random models in paleobiology showed that some random walk in branching and diverging species may produce a pattern of distribution of clades akin to what the fossil records actually display (Fig. 4). From a different viewpoint, McShea (2005) designed simulations to investigate the possible increases in the complexity of organisms across phylogenetic time. McShea

---

[9] I say here "processes", but some of these processes are not "causal processes" stricto sensu, they are, as illustrated, rather causal relations in the sense of counterfactual dependence. What is modeled, through what's called here "pure possible processes" as representational core of the simulation, is thereby not only causal processes stricto sensu but causal relations, in whatever sense of causation can hold among metaphysicians.

challenges here the previous view of the increase of complexity that was held by Bonner (1998)—which hypothesized that there was first selection for large size, and that complexity comes as a consequence of large size, and thus the increase of complexity was a by-product of selection for size. Thus, some of his simulations have no selection in them; instead, they model the diversification of species as a sort of diffusion process, with a fixed boundary.[10] So here, as with the ecological neutral model, computer simulations intend to determine what phylogenetic patterns could be produced by a specific process alone—a process that is somehow a random walk (Fig. 5).

Notice that this claim does not mean that a simulation includes only one single process whereas several of them occur in reality. A simulation may instantiate several processes, for example selection at two distinct levels (genes and organisms for example). The main point mostly concerns the feature of "transparency" proper to simulation: in contrast to real experiments, where several processes can go undetected and be intertwined with the ones the experiment targets, in a simulation only those processes occur that are defined in the algorithm (the undetected processes are just undefined). "Pure" processes exclude the "impurity" constituted by the mixing of well-defined and undefined causal processes. And, in contrast to mathematical models, simulations display the connections between these processes and the patterns they possibly yield.

Therefore, the main epistemic role of computer simulations in evolutionary theory consists in studying those pure possible processes which are likely to yield some pattern. In other words, instead of proving that some process caused some evolutionary result, they provide candidate causal explanations in the following form: "if pattern X is met, then process *x* is likely to have produced it"—the implication being that other causal processes may have been at work but were not significant in regard to this particular outcome.

3 Consequences: the epistemology of computer simulations in evolutionary biology

On this basis, I will now ask what can simulations can do in evolutionary biology—which means: what do they explain and how do they explain it; how do they relate to mathematical models; and how are they validated?

3.1 Explanatory status of evolutionary simulations

Actually, since weak and strong simulations are generally distributed across the classification sketched in Sect. 1.2.2 (i–iii), they cannot require a unified answer to the question of what and how they do explain: strong simulations are often not theory based, weak simulations are theory based (by definition, they rely on the theory that supports the hypothesis they intend to test), and weak simulations may

---

[10] As McShea and Brandon (2011) argue, this diffusion process constitutes the first process in evolution—entailed by the mere fact of variation, that is logically prior to natural selection.

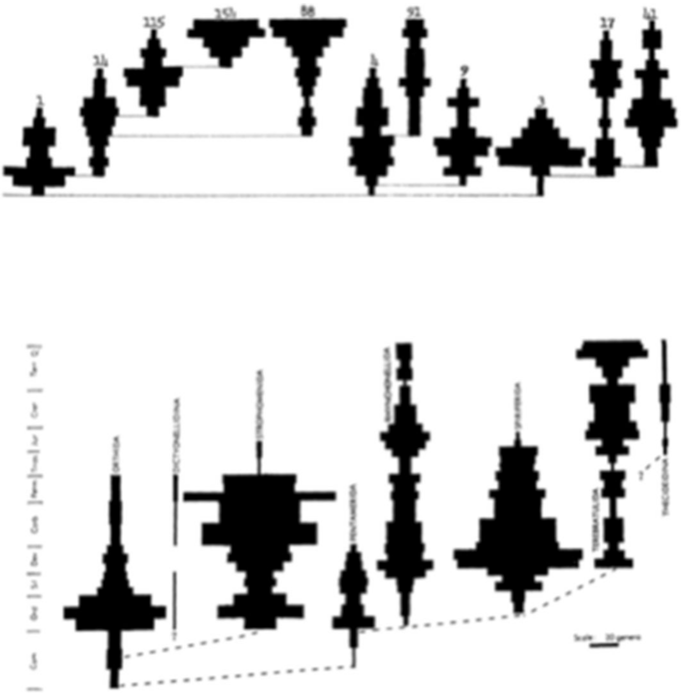

Fig. 4 The similarity between the actual fossil record, order of brachiopods (distribution of clades over times) (*bottom*) and the result of a random simulation. Reproduced from Gould et al. (1977, Fig. 1). Used with permission from the Paleontological Society

be designed in order to test theories whose models are non tractable—which was often the case with the neutral theory in ecology. So, to answer the remaining questions, I will start with the minimal characterization of computer simulations in evolutionary biology demonstrated above: they provide candidate explanations (pure possible processes) for evolutionary patterns. From this, we arrive at two consequences that define the explanatory scope of simulations.

(i) Simulations may *help compare the likelihood of different hypotheses based on distinct pure possible processes in regard to a given pattern*. Recall that the likelihood L of a hypothesis is the probability P of the data conditional on the hypothesis [L = P (data/hypothesis)]. Now, let us suppose n hypotheses ($H_1 \ldots H_n$), each of them claiming that one causal process produces an outcome. For example, the hypotheses might claim that species selection (H_1), or individual selection (H_2), or genic selection (H_3) is responsible for a trait such as territoriality, or for genomic variability. Next, each hypothesis defines a pure possible process instantiable in a weak simulation (e.g. a regime of selection). Thus, each hypothesis allows one to design a specific simulation to instantiate the hypothesized causal process as a pure possible causal process. Now, we can run each simulation a certain amount of times; and in general, the frequency of occurrences of the outcome comparable to the extant data indicates the probability of obtaining such data, were such a causal

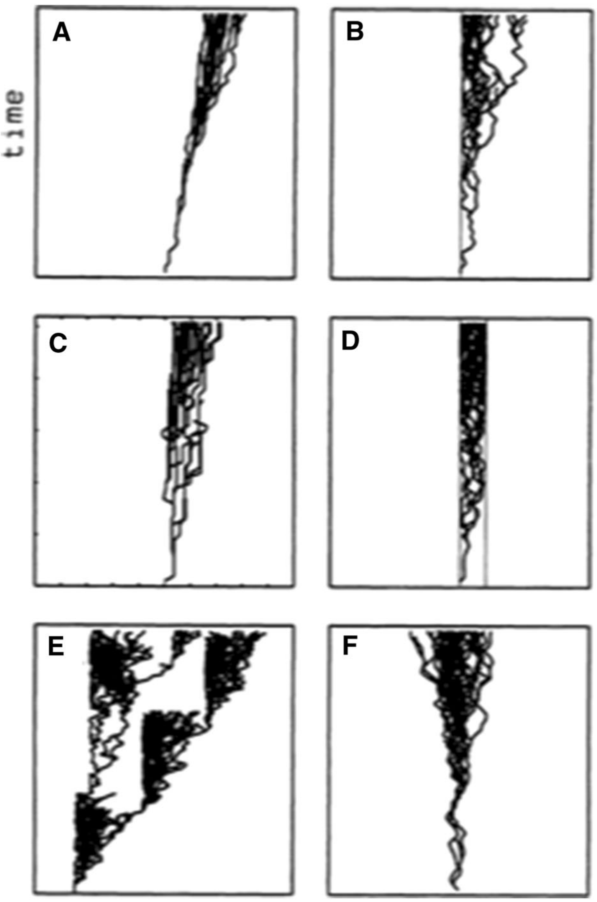

Fig. 5 Various kinds of diversification of a group as outputs from simulations that do not incorporate natural selection (in McShea's view, complexity is defined in terms of number of cell types so diversification is a proxy for complexification). In case B, a lower boundary in complexity explains that there is a trend towards complexification, just by diffusion—this is a passive trend. Reproduced from McShea (1996, Fig. 7). With kind permission from *Evolution*

process to take place alone. Therefore, this gives us the probability of the expected outcome—as a state comparable to the extant data—, conditional on the fact that the hypothesis citing this causal process does hold; that is, such a procedure gives us the likelihood of the hypothesis. This in turn gives us a measure of the plausibility that, the pure processes cited in the most likely hypothesis have in the reality a higher causal impact than the others.

When the simulation is deterministic, the fact that the desired outcome (i.e., the one that is comparable to extant data) is reached significantly more often in the simulations of hypothesis H_i than it is reached in simulations of the alternative hypothesis H_j (in the sense that more initial conditions support it in the case of H_i) means that—in the absence of any knowledge of the initial conditions—the simulated hypothesis H_i is more likely than the alternative hypothesis. Hence, by counting the frequencies of simulation rounds which deterministically yield the expected outcome, or by comparing the frequencies of expected outcomes in the various rounds of a stochastic simulation, we can compare the likelihood of various hypotheses via the simulations, and therefore discriminate among the more or less likely candidate explanations (even though the simulations cannot provide us with the explanation: they just tell us which ones are more or less likely to be correct, among a set of hypotheses that we considered, but the true explanation, in terms of real processes, can always be very different).

(ii) Simulations can *provide null hypotheses for the study of collective behaviours and complex systems*. The case of neutral models in paleobiology by Gould and colleagues exemplified the fact that "pure possible processes" can also be "random processes".[11] In this sense, the simulations also allow one to identify what patterns would likely be yielded by random events. Such a simulation instantiates a null hypothesis (let's call it H_0): to corroborate any alternative hypothesis one should prove that under H_0 the extant data—in the Gould et al. (1977) case, a pattern in the fossil record—is less likely.

The example cited above concerning McShea's (2005) model of the increase of complexity in phylogenetic time provides a null hypothesis for explanations of complexity. If the acknowledged complexity in real data matches the kind of pattern that is produced in the model—by what McShea calls a "passive trend", meaning a directional variation in the absence of selection (see Fig. 5)—then it is less likely that such an increase is due to selection. In other words, these models give us null hypotheses for the explanation of complexity increase.

This is also the case with the neutral models in ecology: some interpretations (e.g. Holt 2006) define the neutral model as the null hypothesis to which alternative explanations considering selection effects ("niche effects", to use the ecologist's language) should be compared. And the main problem for ecologists is that this null hypothesis matches the data quite correctly.[12]

Especially in the case of neutral models in ecology or in evolution, the randomness can be understood as the consequence of certain values of variables being 0 or 1—namely, those variables describing relative fitness.[13] Therefore, the model shows what would happen if there were no selection, and hence it can be understood as a null hypothesis. The neutral ecological model here, for which all fitnesses are equalized (a statement called "ecological equivalence assumption"; see Hubbell 2001), of course instantiates such a situation. But turning to population

---

[11] A neutral model, be it in ecology, paleontology or genetic evolution, models no specific causal process, for example it sets to 0 all parameters on which causal processes depend; however for the epistemological investigation presented in this paper we can call "pure possible process" what is going on in such models, in a way analogous to models of alternate causal hypotheses—"process" here is used in the sense that we talk of "stochastic processes", "chance processes", etc. We leave aside the metaphysical question of whether (some) chance processes are causal processes.

[12] See, for example, Bell et al. (2006, p. 1382): "It was surprising to find that spatial neutral models give rise to frequency distributions of precision that are very similar to those estimated from biological surveys, as a consequence of the spatial patterns produced by local dispersal alone".

[13] If all relative fitnesses in the agents of a model are equal to 1, of course there is no natural selection.

genetics, all models that simulate the stochastic variation of allele frequencies (called "random genetic drift") in terms of diffusion models are also obviously providing null hypotheses for our understanding of evolution. If a hypothesis is that something exists because of selection, then such a hypothesis can be contrasted with the expectations given by this null hypothesis.

To sum up, computer simulations cannot directly prove anything by themselves when it comes to explain how the world really is: the world is defined by "impure processes", and thus simulations are not able to represent accurately what is going on in real biology. However, they provide candidate explanations that are in need of further evidence. To this extent, they could be understood as formulating and testing what Brandon (1990) calls "how possibly explanations"—namely, propositions of causal processes that could possibly have produced the patterns we find out there (see also Dutreuil, this issue). As Brandon notices, "how actually explanations" are scarce in evolutionary biology—given that such explanations would require evidence regarding genetic make-ups, ecological contexts and phylogenies; and we are rarely lucky enough to get all that. Thus simulations, which are capable of providing and assessing "how possibly explanations", are a very frequent tool in evolutionary biology. This feature of computer simulations in evolutionary biology allows us to sketch an answer to the validation problem raised above.

3.2 Validation and Levins triangle

Given the huge differences between kinds of simulations in evolutionary biology, the question of their validation can benefit from being addressed in the context of the Levins triangle, where the distinctions between their proper epistemic values can be made salient. To introduce the question, I start sketching the usual explanatory process that involves weak simulations, and then compare it to an analogous process involving analytical models. Then I'll turn to strong simulations (3.2.2.) and finally situate these teachings in the larger context of the Levins triangle (3.3.).

3.2.1 Case of weak simulations: analogies

The building and use of weak simulations is often the following:

(1) Independent evidence suggests that some processes are possible (for example, that more species form when geographical barriers exist than when they have a definite adaptive potential, or that birds in a flock may follow simple individual rules);
(2) Scientists then design a simulation that implements the main hypothesis in algorithms, which allows them to test the hypothesis that such a pure possible process is indeed likely to produce the intended pattern of data as an outcome (for example, speciations in Maley's model, or the flock's flight behaviour);
(3) If this is the case, more independent evidence has to be gathered in order to confirm that such a "pure possible process" was indeed really occurring (for

> example, more data about geographical barriers and phylogenies, in Maley's case, or more data about the behaviour of actual birds, in the case of flocking boids).

Now, this epistemic process is not wholly special, and should be put in the context of discussions about the methodological steps required in building and validating any model. Important debates about validation revolved around the distinction between verifying (i.e., whether a model correctly computes its solution) and validating (i.e., whether a model is appropriate to its target system). Many of the debates are affected by the question of whether in principle these are two distinct operations, or whether they are intertwined; they are also affected by the fact that the code, as well as the simulation algorithms themselves, has to be verified (Parker 2013). In any case, besides this first validation issue, there is a second one which arises from the basis of the epistemic process I just sketched above.

The question of how we validate simulation models from this viewpoint indeed resembles a classical issue in the epistemology of analytic models, namely, the "calibration vs. confirmation" problem (e.g. Werndl 2013). Quickly said, when we have a set of data and several candidate models to explain this data, we may use two distinct operations with our analytic models: calibration and confirmation—but the same data may be used either to validate or to confirm a model, leading to a kind of underdetermination issue, as I make it clear now. First, recall that two mathematical functions may be very different and yet still yield very close empirical predictions within the domain where we are considering them. For example, a set of data can be compatible with an exponential function such as $y = f(x) = \exp(-ax) + b$, or with a hyperbolic function such as $y = g(x) = 1/cx^{-d}$, or with a specific quadratic form. Granted, they are very dissimilar, but the results $y = F(x)$ are non-discriminable when F is f or g—at least for the range of values we are considering when trying to fit the data (of course, this is only valid with some parameter values—a, b, c, and d in my example). Once a function—f or g—that seems plausible is chosen, one can use the data to calibrate the function (specify the parameters); and then, one will use (possibly other) data to confirm the specified (calibrated) function, that is, to test whether it correctly predicts data sets. Calibrating a model means having already chosen one kind of function (e.g. exponential) and then using the data to fix the parameters in a way that the outcome of the function matches the data; confirming means using data to corroborate the choice of this (exponential) function against another one.

An interesting example of this can be found in what is called "time discounting" in behavioural economics. It has been shown (e.g. Thaler 1981) that when they are asked to make choices at various moments in time (for example, whether they would prefer to have $100 now or $110 in a month), people—rational agents, according to a neoclassical assumption—manifest interesting inconsistencies. Essentially, they tend to show a clear preference for the present (i.e. for the $100 now), which is salient in the high discounting rates they apply to immediate moments in time. But this tendency is inversed when people make the same choice for future delays ($100 in a year or $110 in a year and one month). In this case, they tend to prefer the biggest but most distant alternative. If their choices were consistent through time, they would maintain their preference for the closest reward, no matter the delay. However, a controversy rages about not only the explanation of such discounting, but, before that, its genuine pattern. It is often said that discounting is hyperbolic (Strotz 1956; Laibson 1997), fitting the observed inconsistency. Once this pattern has been determined, then additional data will be used to calibrate the model. However, some people have argued that discounting actually follows an exponential function, first elaborated by Samuelson (1937); in this case, additional data are not used to calibrate the hyperbolic function model, but to confirm an exponential function model independently calibrated. Of course, the causal hypotheses to explain the discounting may be very different if the true discounting function is hyperbolic rather than exponential (Retz Lucci 2013). In any case, this easily shows that the data may be used to either calibrate a hyperbolic function, or to confirm a hyperbolic function against an exponential function (or inversely).

To sum up: if we have independent evidence that the accurate model is an hyperbolic one (perhaps because we have plausible reasons to think that the mechanisms that produce the data are such that they will yield an exponential function), the extant data can be used to calibrate the model. Further data will then be used to confirm that this is a correct model. So here we have a *circular* process: the dual use of data for calibrating and then confirming; in the case of computer simulations, I sketched an explanatory process that was also circular, involving at two stages the appeal to data (steps 1–3 at the beginning of this section). There is an obvious parallel between the circularity in these two protocols. This would lead to the conclusion that nothing really specific characterizes the validation question for simulations: in any case there is the same circular validation process that ties model and data.

However, the circular explanatory process that I just sketched above does not cover cases where one looks for null models. The null models themselves actually enter into the validation process of other models—namely, models of alternative hypotheses. It also does not cover a frequent use of what we called “strong simulations”. Since strong simulations do not implement hypotheses about the world, they can be used neither to discriminate against other hypotheses, nor to forge null models. And since we do not have data against which the hypothesis of strong simulations could be tested, the calibration stage seems meaningless here. For this reason, the validation question addressed above does not really concern strong simulations.

3.2.2 Back to strong simulations

In principle, the two kinds of simulations must be distinguished according to how they deal with pure possible processes. Weak simulations, which we have considered so far in this section, are used to test whether some possible processes can result in forming some type of patterns. On the other hand, I argue that strong simulations are used to explore the potentialities of a *kind* of pure possible processes.

Ray’s Tierra simulations, for example, explore the kind of evolutions that very simple variants and replicators can undergo, independent of the determinate features of the replicators we find in real life (genes, epigenes or “memes”) (Ray 1992). So, when more generally AL practitioners claim that they essentially investigate what evolutionary entities can be and do, this means that their strong simulations explore the outcomes of pure possible processes of selection and variation, in absence of any material constraints on the entities under selection.[14] Such an exploration does not focus on *one* pure possible process, in the sense of a process that could exist in the actual world but is always intertwined with other processes (for instance: gene selection, species selection, etc.); it focuses instead on a *kind* of such processes; natural selection, for instance, with no specification of what the entities under selection are, or what the nature of the inheritance system is. Such processes are likely to be realized in many possible worlds, and therefore intertwined by many other kinds of pure possible processes, some of them not being empirically instantiated in the actual world. The pure possible processes extant and intertwined in the actual world are just instances of such a type or species of process (e.g., gene selection instantiates the type "selection"). Such a focus on kinds or types of pure possible processes manifest in AL characterizes strong simulations in general.

Thus, where weak simulations explore the ability of pure possible processes to yield specified patterns, strong simulations explore the possible *kinds of patterns* produced by different *kinds of processes*. The latter realize internal causal processes that have therefore features not isomorphic to pure possible processes in the world, but to general features of families of processes that can occur in various possible worlds, without necessarily being realized in the actual world.

Yet it appears that the difference seems to be only one of degree: testing a hypothesis about the production of speciation by geographical barriers through a simulation (as Maley did) concerns a process that is quite general—since it is not tied to a specific population; further, because they are not actual processes but what one could call ideal processes, "pure possible processes" are of course quite general. Thus, in regard to the division between strong and weak simulations, one could object that no criterion is given to distinguish between pure possible processes and kinds of processes; in the same way, the distinction between kinds of patterns and patterns seems not to be so robust.

But another important discriminating consideration is the explanatory intent of the simulation: because weak simulations are intended to test a hypothesis, they are supposed to be *compared* against data—extant patterns, and evidence about extant processes. This comparison is involved in the general circular validation process I just described in the previous section. To emphasize the importance of the difference in explanatory intent, think of Reynolds' boids: if one focuses on testing the hypothesis that collective behaviour does not need central control, then they are *weak simulations*; if the aim is to understand all kinds of swarming behaviour available on the basis of simple rules, then they are *strong simulations*.

On the other hand, strong simulations are *exploratory*; therefore, they cannot enter the same circular process of validation by comparison with extant data since there is no data about the different *kinds of patterns* across possible worlds. For

---

[14] These strong simulations have the same epistemic function as what Weber (2014) calls "experimental models"—namely, processes and systems that are designed in order to experiment and test hypotheses about one very general kind of system.

example, no one has uncontroversial data about phylogenetic patterns in evolutionary systems that are more general than DNA inheritance and carbon-based organisms; but data across possible worlds is precisely the data needed to say anything about the "kinds of patterns" that strong AL simulations would look for. Thus, with all criteria (considering kinds of processes, considering kinds of patterns, lack of extant data), one can still distinguish between strong and weak simulations from an explanatory viewpoint. However, this distinction still seems to merge into a continuum; because, as said above, "pure possible processes" (for example, selection) can be ranged along a continuum—from the most individualized (pure genetic selection in a DNA world) to the most general (pure replicator selection, "replicators" being any entity that replicates). I want to emphasize, however, that one can distinguish two poles within this continuum in a non-arbitrary way, which I call strong and weak simulations.

3.3 Mapping simulations on Levin's triangle

In order to understand the various epistemic functions fulfilled by simulations in evolutionary biology it is helpful to turn to Levins triangle and to understand how different simulations trade-off differently between distinct epistemic values. The difference between weak simulations and strong simulations appears here as a difference between realism and generality. Because weak simulations somehow test hypotheses about the causes of real patterns in biology, they may often aim at some form of realism. On the other hand, strong simulations focus on kinds of possible processes (for example, "selection in itself", without regard for the type of selectable entity—not natural selection of individual organisms, of genes, etc.), and to this extent they can be very general.

The two types of simulations thus form two zones within Levins' triangle—one tending towards realism and the other towards generality (Fig. 6). The more one moves towards generality, the less validation protocols such as the one I described above (as analogous to the confirmation/calibration dialectics for mathematical models) are taking place. Instead comes an assessment of another nature—namely, the assessment of the *fruitfulness* of the simulation (see Colyvan 2001, p. 79). Being general, it cannot actually be compared to sets of actual data; but one can instead measure the extent to which it provides insights about the variety of patterns likely to be produced by one kind of process. As one goes towards strong simulations, validation protocols are substituted with fruitfulness assessments. The precise criteria for fruitfulness assessments have to be examined elsewhere.

Of course, the continuum between weak and strong simulations extends between two of the zones in the triangle that I just indicated. For instance, this means that the "pure possible processes" addressed in weak simulations can be seen—once they are taken very abstractly—as "kinds of possible processes" investigated by strong simulations: in this case, you would go smoothly from the pole of weak simulations to the pole of strong simulations.

It should also be noted that the simulations may score differently regarding their precision: some weak simulations will be precise in the sense that they provide quantitative results; while some strong simulations such as Ray's Tierra may be

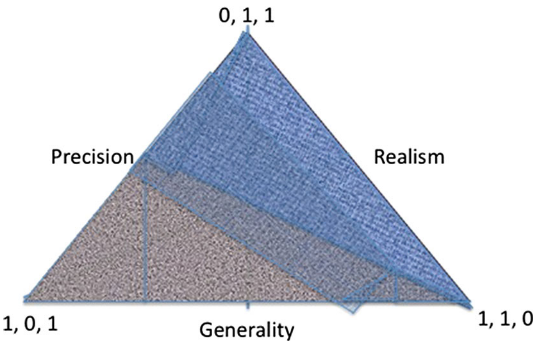

Fig. 6 Strong and weak simulations in the Levins triangle, here divided into two zones, one—*brown*—along the Generality border (*below*) and the other—*blue*—along the Realism border (on the *right*). Of course these two zones would in fact not be sharply separated. This intends to represent the fact that the strong simulations will probably be met in the former zone, and the weak simulations in the weaker zone; and in each zone, the frequency of such simulations increases while one moves toward the border. In the intermediary region, simulations of both kind can be found (drawing by the author)

considered only slightly precise since they offer mere qualitative scenarios. But other strong simulations (such as the Echo ecologies) are also precise since they allow one to measure and quantify the outcomes. Thus, strong simulations occupy a zone in the Levins triangle that is concentrated on the "generality" border, ranging from "quite precise" to "imprecise"; and weak simulations are more centred along the realism border, but also accumulate along the "precision" zone.

The validation problem (as it has been exposed in Sect. 1.2.3) thus concerns simulations when they are understood as belonging to the realism zone inhabited by weak simulations; and in this case, as I said earlier, it is analogous to the usual calibration/confirmation dialectics for analytic models. For this reason, it is plausible that there is no specific common property of simulations in evolutionary biology that would justify giving all of them a specific epistemic status regarding their validation. What is more interesting, rather, is the way they can be distributed along the Levins triangle. It might be that this distribution—when it is explored in detail—will be quite different from the distribution of analytic models in the triangle; as well as, possibly, from the distribution of some models made of real experimental devices.

4 Conclusion

Generally speaking, algorithmic devices only permit us to single out one or a few type of entities. In this sense, they are focusing only on the pure processes involving solely those entities. In evolutionary biology, this means that computer simulations are well suited to explore certain levels of selection, certain regimes of selection, or simply null models for selectionist explanations.

As shown here, one of the intriguing features of computer simulations in evolutionary biology is that they are divided between strong and weak simulations, which fulfil different epistemic aims. This distinction must be taken into account when enquiring about the explanatory scope and status of these simulations, and yields only complex answers—which were sketched in Sect. 3. Even though both types of simulations turn out to be rather evenly distributed across a continuum (than opposed as two types of models), Levins triangle has proven itself to be very useful in building a systematic perspective on the explanatory use of computer simulations in this field. It also constrains the form of the validation problem for computer simulations in evolutionary biology: the more you go towards weak simulations, the more validation problems can be understood in a general circular scheme akin to the calibration and confirmation dynamics in analytic models; the more one becomes interested in strong simulations, the more validation stops being the real issue, especially because "realism" gradually no longer becomes the main epistemic value to fulfil—hence, the more one is compelled to consider other kinds of epistemic assessments, such as fruitfulness, which needs to be investigated in another paper.

Considering computer simulations in evolutionary biology cannot provide a univocal answer to the three questions exposed in Sect. 1. Their variety means that they do not share a common explanatory scope, their validation cannot be understood in a common way, and what they have in common with experiments or mathematical models depends upon whether they are strong or weak simulations; as strong simulations, they resemble some experimental models designed to explore a range of behaviours, such as the bacterial systems of Lenski and Travisano's (1994; cf. Weber 2014); as weak simulations they are chosen and evaluated in the same way as mathematical models are, through a circular process of assessment. However, all these simulations articulate in various ways the two aspects that evolutionary biologist have distinguished and then related ever since Darwin: the patterns (e.g. the Tree of Life) and the processes (e.g. natural selection) of evolution. Because the simulations display in themselves a connection between processes and patterns—as instantiating pure possible processes, and featuring the types of patterns these yield—they actually constitute a rich and original way to investigate this connection. And doing so, these simulations present us, through their variety, with a crucial aspect of the way computer simulations as such modify our capacity to acquire knowledge.

Acknowledgments The author warmly acknowledges Isabelle Drouet and Sébastien Dutreuil for invaluable critiques and suggestions, as well as two anonymous reviewers whose careful reading and constructive criticism considerably improved the paper. Thanks to Adam Hocker for a thorough language checking, and to Staffan Müller-Wille for his precious final suggestions.

References

Adami, C. (2002). What is complexity? *BioEssays, 24*, 1085–1094.

Anderson, P. W. (1972). More is different: Broken symmetry and the nature of the hierarchical structure of science. *Science, 177*, 393–396.

Archetti, M., Scheuring, I., Hoffman, M., Frederickson, M., Pierce, N., & Yu, D. (2011). Economic game theory for mutualism and cooperation. *Ecology Letters, 14*, 1300–1312.

Bedau, M. (2003). Downward causation and the autonomy of weak emergence. *Principia, Revista Internacional de Epistemologia, 6*, 5–50.

Bedau, M. (2008). Is weak emergence just in the mind? *Minds and Machines, 18*, 443–459.

Bedau, M., Snyder, N., & Packard, N. (1998). A classification of long-term evolutionary dynamics. In C. Adami (Ed.), *Artificial Life VI* (pp. 189–198). Cambridge: MIT Press.

Bell, G., Lechowicz, M., & Waterway, M. (2006). The comparative evidence relating to functional and neutral interpretations of biological communities. *Ecology, 87*(6), 1378–1386.

Bonner, J. T. (1998). *The evolution of complexity by means of natural selection*. Princeton: Princeton University Press.

Brandon, R. (1990). *Adaptation and environment*. Cambridge, MA: MIT Press.

Colyvan, M. (2001). *The indispensability of mathematics*. New York: Oxford University Press.

Crubelier, B., Preux, P., & Cambier, C. (1997). Studying adaptation with Echo. ftp://ftp.cogs.susx.ac.uk/pub/ecal97/online/F045.ps.gz.

Damuth, J., & Heisler, I. L. (1988). Alternative formulations of multilevel selection. *Biology and Philosophy, 3*, 407–430.

Dennett, D. (1995). *Darwin's dangerous idea*. New York: Simon & Schuster.

Eldredge, N. (1985). *Unfinished synthesis: Biological hierarchies and modern evolutionary thought*. New York: Oxford University Press.

Epstein, J. (1999). Agent-based computational models and generative social science. *Complexity, 4*(5), 41–57.

Epstein, J. (2002). Modeling civil violence. *Proceedings of the National Academy of Sciences of the United States of America, 99*(3), 7243–7250.

Frigg, R., & Reiss, J. (2009). The philosophy of simulation: Hot new issues or same old stew. *Synthese, 169*, 593–613.

Gaucherel, C., Munoz, F., & Bérard, S. (2011). Equation against algorithm: Differences and choosing between them? *Acta Biotheoretica, 59*(1), 67–79.

Gillespie, J. (2004). *Population genetics*. New York: Oxford University Press.

Gould, S. J., & Lloyd, E. A. (1999). Individuality and adaptation across levels of selection: How shall we name and generalize the unit of Darwinism? *Proceedings of the National Academy of Sciences of the United States of America, 96*, 11904–11909.

Gould, S. J., Raup, D., Sepkoski, J., Schopf, T., & Simberloff, D. (1973). Stochastic models of phylogeny and the evolution of diversity. *The Journal of Geology, 81*(5), 525–542.

Gould, S. J., Raup, D., Sepkoski, J., Schopf, T., & Simberloff, D. (1977). The shape of evolution: A comparison of real and random clades. *Paleobiology, 3*, 23–40.

Holland, J. (1995). *Hidden order: How adaptation builds complexity*. Readings: Helix.

Holt, R. D. (2006). Emergent neutrality. *Trends in Ecology & Evolution, 21*(10), 531–533.

Hraber, F. S., & Jones, T. (1994). Modeling complex adaptive systems with Echo. In R. J. Stonier & X. H. Yu (Eds.), *Complex systems, mechanisms of adaptation* (pp. 3–21). Amsterdam: IOS Press.

Hubbell, S. P. (2001). *The unified neutral theory of biodiversity and biogeography*. Princeton: Princeton University Press.

Hughes, R. (1999). The Ising model, computer simulation, and universal physics. In M. Morgan & M. Morrison (Eds.), *Models as mediators* (pp. 97–145). Cambridge: Cambridge University Press.

Humphreys, P. (2004). *Extending ourselves*. Oxford: Oxford University Press.

Huneman, P. (2008a). Emergence and adaptation. *Minds and Machines, 18*, 493–520.

Huneman, P. (2008b). Combinatorial vs. computational emergence: Emergence made ontological? *Philosophy of Science*, *75*, 595–607.

Huneman, P. (2012a). Determinism and predictability and open-ended evolution: Lessons from computational emergence. *Synthese, 185*(2), 195–214.

Huneman, P. (2012b). Computer sciences meet evolutionary biology: Issues in gradualism. In J. L. Torres, O. Pombo, J. Symons, & S. Rahman (Eds.), *Special sciences and the unity of science* (pp. 200–225). Dordrecht: Springer.

Huss, J. (2008). The shape of evolution: The MBL model and clade shape. In D. Sepkoski & M. Ruse (Eds.), *The paleobiological revolution* (pp. 339–358). Chicago: University of Chicago Press.

Kettlewell, H. B. D. (1955). Selection experiments in industrial melanism in Lepidoptera. *Heredity*, *9*(3), 323–342.

Küppers, G., & Lenhard, J. (2005). Validation of simulation: Patterns in the social and natural sciences. *Journal of Artificial Societies and Social Simulation, 8*(4). http://jasss.soc.surrey.ac.uk/8/4/3.html.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics, 112*, 443–477.

Lenski, R., & Travisano, M. (1994). Dynamics of adaptation and diversification: a 10,000-generation experiment with bacterial populations. *Proceedings of the National Academy of Sciences of the United States of America, 91*(15), 6808–6814.

Lethinen, A., & Kurikovski, J. (2007). Computing the perfect model: Why do economists shun simulation? *Philosophy of Science, 74*, 304–329.

Levins, R. (1966). The strategy of model building in population biology. *American Scientist, 54*, 421–431.

Lewontin, R. C. (1970). The units of selection. *Annual Review Ecology Systematics, 1*, 1–18.

Lewontin, R., & Dunn, L. (1960). The evolutionary dynamics of a polymorphism in the house mouse. *Genetics, 45*, 65–72.

Lopez, U., Gautrais, G., Couzin, D., & Theraulaz, G. (2012). From behavioral analyses to models of collective motion in fish schools. *Proceedings of the Royal Society Interface Focus, 2*, 693–707.

Maley, C. (1998). Comparing causal factors in the diversification of species. PhD Dissertation, MIT, accessible at *InterJournal Complex Systems,* 241. http://www.interjournal.org/manuscript_abstract.php?28997.

Margolus, D., & Toffoli, L. (1987). *Cellular automata*. Cambridge: MIT Press.

Maron, M. (2004). Evolution of industrial melanism: A spatial, predator-prey genetic algorithm. http://www.brainoff.com/easy/moth/report.pdf.

Matthewson, J., & Weisberg, M. (2009). The structure of trade-offs in model building. *Synthese, 170*(1), 169–190.

McShea, D. (1994). Mechanisms of large-scale evolutionary trends. *Evolution, 48*(6), 1747–1763.

McShea, D. (1996). Metazoan complexity and evolution: Is there a trend? *Evolution, 50*(2), 477–492.

McShea, D. (2005). The evolution of complexity without natural selection: A possible large-scale trend of the fourth kind. *Paleobiology*, *31*(2), 146–156.

McShea, D., & Brandon, R. (2011). *Biology's first law*. Chicago: University of Chicago Press.

Miller, G., & Todd, P. (1995). The role of mate choice in biocomputation; sexual selection as a process of search, optimization and diversification. In W. Banzhaf & F. Eckmann (Eds.), *Evolution and biocomputation: Computational models of evolution* (pp. 169–204). Dordrecht: Springer.

Morrison, M. (2009). Models, measurement and computer simulation: The changing face of experimentation. *Philosophical Studies, 143*, 33–57.

Norton, S., & Suppe, F. (2001). Why atmospheric modelling is good science. In C. Miller & P. Edwards (Eds.), *Changing the atmosphere: Expert knowledge and environmental governance* (pp. 88–133). Cambridge: MIT Press.

Okasha, S. (2006). *Evolution and the levels of selection*. New York: Oxford University Press.

Parker, W. (2013). Computer simulation. In S. Psillos & M. Curd (Eds.), *The Routledge companion to philosophy of science*. London: Routledge.

Peck, S. L. (2004). Simulation as experiment: A philosophical reassessment for biological modeling. *Trends in Ecology & Evolution, 19*, 530–534.

Peschard, I. (2010). Modelling and experimenting. In P. Humphreys & C. Imbert (Eds.), *Models, simulations, and representations* (pp. 42–61). London: Routledge.

Pocheville, A. (2014). Ecological niche. In T. Heams, P. Huneman, G. Lecointre, & M. Silberstein (Eds.), *Handbook of evolutionary thinking in the sciences*. Dordrecht: Springer (in press).

Rasmussen, S., & Barrett, C. L. (1995). Elements of a theory of simulation. In F. Moran (Ed.), *Advances in artificial life: Third European conference on artificial life* (pp. 515–529). Dordrecht: Springer.

Ray, T. (1992). An approach to the synthesis of life. In C. Langton (Ed.), *Artificial Life II* (pp. 371–408). Boston: Addison Wesley.

Reich, E. S. (2012). Flaws found in faster-than-light neutrino measurement. *Nature News*. doi:10.1038/nature.2012.10099.

Retz Lucci, C. (2013). Time, self, and intertemporal choice. *Frontiers in Neuroscience, Decision Neuroscience*. doi:10.3389/fnins.2013.00040.

Reynolds, C. (1987). Flocks, herds and schools: A distributed behavioural model. *Computer Graphics, 21*(4), 25–34.

Samuelson, P. A. (1937). Note on measurement of utility. *Review of Economic Studies, 430*(4), 155–161.

Strotz, R. (1956). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies, 23*, 165–180.

Taylor, T. (2012). Exploring the concept of open-ended evolution. In C. Adami, D. Bryson, C. Ofria, & R. Pennock (Eds.), *Proceedings of the 13th international conference on artificial life* (pp. 540–541). Cambridge: MIT Press.

Thaler, R. (1981). Some empirical evidence on dynamic inconsistency. *Economics Letters, 439*(8), 201–207.

Toner, J., & Tu, Y. (1998). Flocks, herds, and schools: A quantitative theory of flocking. *Physical Review E, 58*, 4828–4858.

Vicsek, T., Czirok, A., Ben-Jacob, E., Cohen, I., & Shochet, O. (1995). Novel type of phase transition in a system of self-driven particles. *Physical Review Letters, 75*, 1226–1229.

Volkov, I., Banavar, J. R., Hubbell, S., & Maritan, A. (2003). Neutral theory and relative species abundance in ecology. *Nature, 424*, 1035–1037.

Weber, M. (2014). Experimental modeling in biology: In Vivo representation and Stand-ins as modeling strategies. *Philosophy of Science* (in press).

Werndl, C. (2013). Climate models, confirmation and calibration. *The British Journal for the Philosophy of Science, 64*, 609–635.

Winsberg, E. (2010). *Science in the age of computer simulation*. Chicago: University of Chicago Press.

Wolfram, S. (1984). Universality and complexity in cellular automata. *Physica D: Nonlinear Phenomena, 10*, 1–35.