2011 Eastern Division invited paper in the session:
Author Meets Critics: Eric Winsberg, *Science in the Age of Computer Simulation*

Computer simulation and the quest for novel epistemic novelty

Isabelle Peschard

San Francisco State University

At the very beginning of *Science in the age of computer simulation,* Eric Winsberg diagnoses the lack of philosophical interest for simulations, at least until recently, as an effect of the presumption that simulations raise only technical issues because they simply consist in applying theories and revealing the knowledge that they somehow already contain.

By contrast, an overarching theme of the book, is that simulations involve application of theories in a complex and creative way resulting in what deserves to be called "genuinely novel knowledge": "simulation is in fact a deeply creative source of scientific knowledge".

I am intuitively sympathetic to the idea that simulations produce 'genuinely novel knowledge'. But when I think about it, I find myself unsure about exactly what that means. Simulations produce knowledge, they produce information that was not already there, and so is novel. But that is too obvious to be what Winsberg has in mind. There must be something more that the 'genuinely' is pointing at. The question is 'What?'.

So what I would like to do in this discussion is to try to clarify, by going through different chapters, the sense in which the results of simulation may be said to constitute *genuinely novel* knowledge. A good way to understand what this novelty is is to understand where it comes from- so what I will try to do is to locate, among the different aspects that are discussed in the book of a simulation as a complex process, the source of this epistemic novelty which, according to Winsberg, characterizes simulation results. What aspect(s) of the simulation as a complex process makes it possible?

I. One possible answer is that simulation involves modeling, and modeling is generally not just deriving from theories.

Theoretical models, as has been by now clearly demonstrated (e.g. Morrison, Cartwright et al. 1995, Boumans) are constructed in the sense that they integrate, articulate, combine, in a way that is not necessarily dictated by any theory, having elements that may come from different theories or not come from a theory at all. In this sense, models are a source of epistemic novelty: their epistemic content is not just part of the epistemic content of any old theory. Winsberg does emphasize and comment on the modeling aspect of simulation. But if that were the source of epistemic novelty that he has in mind for simulation, then this novelty, it seems, would not be so novel after all. The claim of novelty would rather look like a reiteration, certainly with new evidence, of the claim made about models and modeling in general.

One may object that there is actually something specific to the modeling involved in simulation. There are some specific constraints that bear on this sort of modeling: in particular, the model that is produced for simulation must be suitable for being transformed into a model that can be, ultimately, implemented on a computer.

It is hard to see, however, how that could justify that one speak of simulation as producing a sort of knowledge that is novel in a way essentially different from the way in which modeling in general produces novel knowledge. It certainly makes it clear that the construction of theoretical models, in the context of simulation, may have to integrate constraints that are clearly independent from the theory, stemming from the implementation on a computer and the specific task of the simulation. It would provide a compelling basis to the claim that models involved in simulation may be used to produce novel knowledge, in the sense of knowledge not implicitly present in the theory. But it doesn't show how there can be a form of epistemic novelty specific to the results produced by simulations.

II. Another possible source of epistemic novelty might be the next step in the simulation process: the transformation of the theoretical model into a model that can be implemented on the computer, a simulation model.

This step can be seen as an additional modeling. Instead of being the construction of a model on the basis of some theoretical principles, it consists in the construction of a simulation model on the basis of the theoretical model, itself constructed on the basis of theoretical principles.

To say that this step is the source of epistemic novelty that distinguishes simulation from mere modeling would seem to imply that there is some epistemic content in the simulation model that was not already in the theoretical model. It suggests that we may think of the simulation model as autonomous with respect to the theoretical model in the way we may think of theoretical models as autonomous with respect to theories/theoretical principles.

The idea of such autonomy is counter intuitive. We think of the simulation model as an implementable form of the theoretical model; different but making sense only relatively to the theoretical model and such that its content will be as good an approximation of the theoretical model as possible.

But to think that way is to make a mistake, according to Winsberg, about the real aim of the simulation: not to produce solutions to the original equations but to have simulation results that are empirically informative. Hence the construction of the simulation model suitable for computation given certain constraints of time, computational power, and accuracy may involve approximations and the addition of terms with no theoretical ground and no other function than to make up for these approximations. One example discussed by Winsberg is the incorporation of a term of so called 'eddy viscosity' in astrophysical simulations. This term produces on the large scale of the computational grid the accumulated effect that real viscosity produces at smaller scales[1].

It is not clear, however, that we can really speak of autonomy of the simulation model in the same sense as the autonomy of theoretical models with respect to theoretical principles. One reason is that theoretical principles are not about any system in particular- as Ronald Giere puts it, they are not representational, by contrast to the theoretical model that is constructed on the basis of these principles. The theoretical model is about a specific system in a way the theoretical principles were not. One might even think that this is where the epistemic novelty generated by the theoretical model comes from: the theoretical model is about a specific system in a way the theoretical principles could not be.

By contrast, the theoretical model and the simulation model are both about a system in particular, and they are about the same system. So in what sense could there be some epistemic content that pertains to the simulation model but not to the theoretical model?

One answer might be that they have different epistemic content in virtue of the fact that the epistemic content of the simulation model is an approximation of the epistemic content of the theoretical model. Simulations make it possible to deal with models that are analytically intractable. And in this sense, as producing approximate solutions for these models, they may be said to enlarge the domain of actual knowledge. But it is not clear that one should then talk of a 'genuinely' novel knowledge produced by the simulation. That we obtain 'only' an approximation could be seen as an epistemic limitation rather than the mark of an epistemic autonomy. Probably, this is why, as we saw, Winsberg insisted that to think of simulation result merely as an approximation is not the right way to think of simulation; it makes the simulation model a mere epistemic shadow of the theoretical model.

What else then may explain the genuine epistemic novelty of simulation? As we saw earlier, and it is true at the level of the simulation model too, the simulationist may have to deploy some techniques of modeling that are novel, specifically tailored for the specific problem of implementation. But again, it is not clear how different it is, epistemically speaking, from constructing a theoretical model that is simple enough, by using drastic approximations, to be analytically tractable. To produce such theoretical models also requires creativity, ingenuity. It may even also appeal to

[1] For another demonstration of the autonomy of the simulation model with regard to theoretical models, see Lenhard REF.

approximations and fictional terms. So what is really novel about simulation epistemic novelty?

III. The next step in simulation to consider as possible source of epistemic novelty is the implementation of the simulation model via an algorithm.

One reason to think that the implementation of the simulation model may be what really distinguishes simulation from mere modeling is that this implementation makes simulation a form of experimentation. And one reason to think that this experimental aspect of simulation might be a key to identifying the source of the epistemic novelty of simulation results is that Winsberg makes a point to argue that "experiments are [not] **epistemically privileged** relative to simulations".

The idea, here, seems to be that in virtue of some similarity between simulation and experimentation, the result produced by simulation are no less epistemically novel than the results produced by experimentation.

Winsberg succeeds in revealing compelling similarity between the simulationist's and experimenter's *methodology* for providing credentials for their results: the simulationist will check her results against other trusted results, theoretical or empirical, and rely on previous models, skills and techniques selected by their previous success.

But this methodological similarity is not sufficient to justify the conclusion about epistemic novelty. What needs to be compared, after all, is what can be learned about a given system.

It is true that beyond the methodological similarity, simulation can be analyzed in three sub-processes that characterize experimental procedure:

1) The preparation of a system (preparation and implementation of the simulation model)
2) The evolution of the system: the autonomous transformation over time of the physical system that implements the computation (Humphreys 1994; Norton and Suppe 2001). It may be argued that, in a certain sense, the model that is simulated also undergoes a transformation (Krohs 2008)
3) The recording, organization, and classification of the results in the form of models of the data (Winsberg 2003).

But one may want to resist the conclusion that simulation is epistemically on a par with experimentation on the intuitive basis that, in experimentation, the system investigated, or target system, is experimented on (manipulated and probed), via instruments used to perform these actions and to record their conditions and consequences. In simulation, by contrast, what is manipulated is only a model of the target system. How good, however, is this intuition? Is it really the case that in experimentation the system investigated is manipulated? One way one may try to undermine this intuition is by examining the function of models both in experimentation and in simulation. And Winsberg does underline the role of modeling in both experimentation and simulation, especially at the level of the construction of models of the data. This is not however the main basis of his argument against the intuition that in experimentation, by contrast to simulation, the target system is manipulated and probed.

Rather, according to Winsberg, both experimentation and simulation consist in the manipulation of an object that stands in for the target system. It will be a physical system in experimentation, a model in simulation. But in both cases, information about the system of interest is obtained indirectly, as the product of an inference from the results of the manipulation.

Winsberg doesn't deny that there are some differences between simulation and experimentation. But he sees these differences as being strictly epistemological, a difference in the kind of justification supporting the inference from the result of the manipulation to claims about the target system, in what serves as basis for the reliability of the inference.

With simulation, it is trust in the simulation model and the different elements involved in its construction (theoretical principles, tricks and physical intuitions). With experimentation, it will be reasons to regard the object manipulated and the system of interest as being, in relevant ways, the same kind of material system. But given that one basis is not, as a matter of principle, more secure than the other, "experiments", Winsberg concludes, "are [not] **epistemically privileged** relative to simulations". (p.70)

Winsberg makes a compelling argument to the effect that *if* in experimentation, just as in simulation, what is manipulated is a system standing in for the target system, then there is no basis for drawing a principled distinction between the epistemic functions of simulation and experimentation. But the premise of the argument, that in both cases what is manipulated is a system standing in for the target system, is questionable. I will discuss two ways in which it might be undermined:

First, it does not seem to be necessarily the case in experimentation, by contrast with simulation, that the system manipulated is different from the target system. And secondly, when the two are distinct in an experiment, the relation between them is different from what it is in the case of a simulation.

- Regarding the first point, the distinction between system manipulated and target system finds its or one of its strongest motivations in the observation that what we want to learn may not be accessible to manipulation (too complex, or already in the past, or still in the future…). For instance, we want to learn about human reactions to drugs, but we manipulate rats. The system manipulated, it then seems, is not the one we really want to learn about. On the basis of this distinction, two different problems arise: one related to the validity of inferences about the system manipulated (problem of internal validity); the other related to the validity of inferences about the target system on the basis of the results about the system manipulated (problem of external validity).

Interestingly, however, Francesco Guala (2008) notes and I believe he is right, that "experimental physicists do not recognize external validity as a separate problem of inference", and more generally, that "experimenters are often concerned with proving the existence of certain mechanisms or phenomena in the lab only, and leave it to policy-makers or applied economists to apply such knowledge in the field". But if experimenters are not themselves concerned with drawing inferences about the system 'in the field', then it is certainly inappropriate to take this system to be the target system of the experimentation. In the simulation, what is identified as target system is the system which the simulation is designed to produce information about (the system represented by the model manipulated). If we are coherent and identify the target system of experimentation in the same way, the system in the field then is not the target system of experimentation.

The system in the field may be the **epistemic motivation**. It motivates the experimental procedure and the epistemic function of this motivation is important. Still, going back to the experimental study on rats, the motivation for such a study is, we said, human reaction to the drug. But it may have been something else. So what system experimental results are about and what system they may be used to learn about are two very different things.

Just as in simulation, in experimentation, the target system should be what the procedure of manipulation is specifically designed to learn about, what the conclusions of the experimental study are about, which is different from the speculations they may encourage. And the results of the experimental study are about the system that is manipulated; at least it is so in a large number of cases, and especially in physics.

- But there is also a class of cases of experimentation where object manipulated and target system do not coincide: typically, when the system manipulated is a sample from a population about which the manipulation is designed to make conclusions.

Under the assumption that the sample is representative of the whole, the results are about the whole population. In this case, the former seems to qualify as much as the latter for the status of target system. And it seems that being representative amounts to standing in, in the way that in a simulation the model stands in for the system it is a model of. There is an important difference though: the model is not *representative* of the system it represents, it is meant to be a *representation* of it.

Mary Morgan (2003) already pointed out this distinction between being a representation and being representative. The distinction is between two different ways in which a system may stand in for another: by contrast to being a representation, the system that is representative of another is only different from it in the way that a part is different from the whole. Morgan speaks of the representative as being 'of the same stuff' as what it is representative of. 'Being of the same stuff' should be taken literally: 'being a part of'.

Wendy Parker (2009) made clear that material similarities between the system manipulated and the target system are not necessarily more informative than the formal similarities relied upon in simulation. But I think that what is important in the idea of 'being of same stuff' is not the idea of a different type of similarities; it is the idea that what is manipulated is not a different sort of system. It is more like a sub-system.

Winsberg 's analysis of the epistemological consequences of the difference between manipulating a model and manipulating an experimental system still applies: the justifications for the results will be of different sorts. But that doesn't justify its *epistemic* conclusion. For there will be other consequences of the difference between manipulating a *representation* and manipulating a *representative* of the target system. And it seems very reasonable to think that these consequences are epistemic. They have to do with what we can learn about the target system: the idea of a representative is the idea of a system that, because it is made of 'the same stuff', will react to the manipulation in the same way as the target system would if it was manipulated.

If all that is exact, it seems not to be the case that, in experimentation, we generally learn about a given target system by manipulating a different system that stands in for the target system in the same way as a model stands in for the target system in simulation. So this idea, that what is manipulated is a system that stands in for the target system in the same way in experimentation and in simulation, cannot justify the claim that experimentation had no epistemic privilege with respect to simulation. And correlatively, it doesn't explain the novel epistemic novelty of simulation.

If there is some genuine epistemic novelty to simulation results it is not in virtue of the procedural similarity with experimentation that Winsberg thinks there is.

But that doesn't mean that the experimental aspect of simulation plays no role in explaining the epistemic novelty of simulation. It just means that, if it does, it is in a different way. But we are back at the beginning: In what way then?

CONCLUSION

Unfortunately I do not have a clear answer. But I have the impression that the reason why we were not able to identify the source of epistemic novelty is that we were looking for the answer in some similarity between simulation and modeling and between simulation and experimentation. If what we want is a source of epistemic novelty shouldn't we rather look specifically at the dissimilarity between simulation and modeling and between simulation and experimentation:

On the one hand, it is true that simulation involves model building: but, by contrast with modeling, it involves the *manipulation* of the model.

On the other hand, it is true that simulation is a form of experimentation: but by contrast with experimentation, it involves the manipulation of a *model* (of the target system).

Winsberg at some point mentions that simulations have a life of their own, as experimental systems do, because simulation techniques are being developed, evolve and travel from one investigation to the other. But there is another way in which the simulation has a life of its own, referring now to the simulation process itself and its temporal evolution, to the way the model unfolds in time. Maybe what matters is that the manipulation of the model enables something to be 'said' about the target system that could not be said by analyzing the model or measuring the state of the target system, something that has to do with the historical development of this state.

Interestingly, Winsberg acknowledges but sets aside the focus of dissimilarities and the views of simulation as lying between theory and experiment without being either. It is not really illuminating for his epistemological project, he says. But what about the epistemic overarching thesis? Could that be that the epistemological project and the epistemic thesis cry for opposite philosophical strategies: examining the similarities of simulation with both modeling and experimenting, for the epistemological project, versus teasing out the dissimilarities to find support for the epistemic thesis?