The Epistemic Superiority of Experiment to Simulation[1]

If we are thinking within our system, then it is certain that no one has ever been on the moon. Not merely is nothing of the sort ever seriously reported to us by reasonable people, but our whole system of physics forbids us to believe it. For this demands answers to the questions “How did he overcome the force of gravity?” “How could he live without an atmosphere?” and a thousand others which could not be answered. – Ludwig Wittgenstein, c. 1950 (published 1969)

Somewhere something incredible is waiting to be known. – Carl Sagan

1. INTRODUCTION

This paper defends the naïve thesis that the method of experiment is epistemically superior to the method of simulation, other things equal, a view that has been rejected by some philosophers writing about simulation, and whose grounds have been hard to pin down by its defenders. There are three challenges I take on in defending this thesis. One is to say how “other things equal” can be defined, another to identify and explain the source of the epistemic advantage of experiment in a hypothetical comparison so defined. Finally, I explain why this theoretical point matters since practical constraints like feasibility and morality mean that scientists do not often face an other-things-equal comparison when they choose between experiment and computer simulation (hereafter “simulation”).

2. TWO METHODS

There is a persistent intuition that experiments are more direct than simulations, that they are in a more direct relationship to the object of study, the material world. This intuition surely has a role in the concern that has been expressed among some experimental physicists that

[1] I would like to thank XXX for helpful discussions of this material.

simulations will come to be preferred because they are generally cheaper, and that this will inhibit discovery of new facts about the world (Humphreys 2004, 133-134). After all, a simulation can only reveal the consequences of knowledge we already possess, or, as Herbert Simon put it, “a simulation is no better than the assumptions built into it” (Simon, 1969, 18).

This latter intuition is supported by an idea that a computer simulation is merely calculating the consequences of a set of theoretical assumptions. There would thus be no information in the results of the simulation that were not already present in the theoretical assumptions – garbage in, garbage out. However there are two problems with this idea. First, the fact that information is present in theoretical assumptions does not mean that we know it is, so even if a simulation is merely calculating, it is giving us new knowledge. One could protest that the epistemic status of the new beliefs we form is no better than that of the theory since it has no grounds independent of our trust in the theory. However, in fact a computer simulation’s results do have sources independent of the theory, since it is a long road from theoretical assumptions to a “solver,” the computer program that will give the final results (Morgan 2002, 2003, 2005; Humphreys 2004; Winsberg 2010). The solver is the product of theoretical ideas, approximation, replacement, and computational lock-in, ingenuity, and necessity. It has “a life of its own” and an independent epistemic status; its credentials are evaluated independently of the theory and it is typically the solver, not the theory, that is revised when predictions resulting from a simulation face recalcitrant experience. This independence of evaluation is the second problem with viewing a simulation as merely calculating the consequences of a theory.

If a simulation has enough distance from the theoretical assumptions to give it an independent status, perhaps the difference between simulation and experiment can be located in the former having too much distance from the target system, the part of the world that the two methods can be used to instruct us about. On one such view, whereas in an old-fashioned experiment one is “controlling the actual object of interest, …, in a simulation one is experimenting with a model rather than the phenomenon itself.” (Gilbert and Troitzsch 1999) An immediate problem with this view is that the actual object of interest of a scientist conducting an old-fashioned experiment to test a hypothesis often extends beyond the sample that can be manipulated in the lab. Ernest Rutherford wanted to use an experiment on a gold sample to draw a conclusion about the structure of all atoms, not only those in other samples of gold, but also in lead, and hydrogen, and phosphorus. A model may be thought to be farther from the world insofar as in order to draw conclusions about the world from it one must make assumptions as to its similarities to the target system. But as evident to us as it may be, the claims that this gold is similar to all other gold and to lead, and hydrogen, and phosphorus, in the relevant respects are also assumptions that must have been justified if the results on this gold are to be generalized to all atoms. In both simulations and experiments the object acted upon is separated from the world to be learned about by a layer of assumptions of relevant similarity (Parker 2005, Winsberg 2009, 2010).

Another way of attempting to make out the intuition that experiment is more direct than simulation focuses on the kind of similarity that obtains between study system and target system. A gold sample is materially similar to all other gold, and to lead, hydrogen, and phosphorus, in the relevant respects, whereas a computer model is similar to the target system only in virtue of its form (Morgan 2002, 2003, 2005; Guala 2005). We can surely grant that there is a continuum with material, associated with experiment, on one end, and formal, associated with (computer) simulation, on the other. And though a computer model may have a similarity of functional organization to all gold and other atoms, and the program's instantiation in hardware may follow relevantly similar dynamical rules of evolution, no hardware we currently imagine will ever be as materially similar to all gold as a sample of gold is. If we thought it were, we would likely think twice about calling the system a simulation.

But while a material/formal, experiment/simulation continuum tracks something in our usage of the words, the purpose of these two methods is to learn things about the world, so the challenge for this view is to explain why the distinction between material and formal similarity makes any principled epistemological difference. As I have said, a claim of material similarity does not come for free but depends on background knowledge or assumptions. Moreover, even if we grant that material similarity will bring strictly more similarity than formal similarity will, because relevant material similarity will include formal or dynamical similarity, one must still say why that additional similarity makes an epistemic difference.

Several authors claim that material similarity is always relevant:

> We are more justified in claiming to learn something about the world from the experiment because the world and experiment share the same stuff. In contrast, inference from the model experiment is much more difficult as the materials are not the same – there is no shared ontology, and so the epistemological power is weaker. (Morgan 2005, 323; cf. Guala 2002, 2005; Harré 2003, 27-8.)

but the claim that "ontological equivalence provides epistemological power" (Morgan 2005, 326) gives no guidance as to why the extra similarity always matters to justification. The proposal is that material similarity makes it easier to justify any claim of relevant similarity. However, this puts the cart before the horse, for the claim of material similarity itself requires justification. Material similarity does not give a metric for epistemically relevant similarity between the study and target systems because establishing such similarity itself adds a layer of epistemic distance, which could in principle be of any length, and material similarity is not necessarily easier to establish than formal similarity.

However, there is a sense in which similarity of material can let the experimentalist get away with less than the simulationist, that is well explained by Francesco Guala (2005), and a version of which is endorsed by Eric Winsberg (2010, 64-69). Both experimentalist and simulationist must insure that their study systems are dynamically similar to the target system. The simulationist does this by making in her study system a model of the dynamics of the target system. The experimentalist can circumvent the need to make a model of the dynamics of the target system if she has reason to believe the study sample and target system are materially similar, because dynamical similarity can be inferred from that material similarity. Rutherford could suppose that this gold *behaves* like all other gold in the relevant respects – whatever they are – because they all *are* gold. Of course that claim of material similarity must be justified, but the simulationist must go further, to make specific commitments about what the dynamics of the target system are – commitments that can be avoided and black-boxed in an experiment – and the simulationist will have no material similarity at all to appeal to. Thus the simulationist seems to be strictly further out on a limb. This contrast plays a role in what I will argue is the superiority of experiment to simulation, although I will not take the contrast to rest essentially on the material vs. formal distinction.

3. OTHER THINGS EQUAL

Margaret Morrison, Wendy Parker, and Eric Winsberg (Morrison 2009, Parker 2009, Winsberg 2009) have denied that the difference between material and formal similarity has epistemic significance per se, and that experiment is a superior method. In partial support of these claims, Parker and Winsberg point out that some simulations are better than the experiments that we are able to do in pursuit of the same question, despite the fact that the experiments would be much more materially similar, for example, same-stuff models of weather and same-stuff models of black holes (Parker 2009, 492; Winsberg 2010, 61). However, this point is not probative for two reasons. One is the qualification to experiments that we are able to do. That there are questions for which the simulation we are able to do is more reliable than any experiment we can do gives no reason to deny the superiority of a comparable experiment that we cannot do. Methods can often be compared even when they cannot be carried out, and the superiority claimed here is epistemic, not pragmatic.

Secondly, in order to isolate the difference that being an experiment or a simulation makes we must compare the two methods *other things equal*. Parker recognizes this but despairs of defining this phrase in the current context since it seems impossible to make the "“same” intervention or make the “same” observations in two experiments in which the systems being intervened on and observed are quite different in material and structure" (Parker 2009, 492). However, the equality needed is not material or structural; it is epistemic. And equality is needed only for those properties that do not distinguish an experiment and a simulation. My procedure for setting up an other-things-equal comparison will be first to identify as many general similarities as possible between the two methods and thereby close in on the differences by elimination. Then, the difference that will be the basis for an argument for the superiority of experiment will emerge when we take an actual experiment and constructively imagine the best possible simulation for addressing the same question, that is similar to the experiment in every possible epistemic way.

In the other-things-equal comparison, the two studies must be aiming to answer the same question. Even beyond this, for all their material and structural differences, the methods of experiment and simulation are remarkably similar epistemically, in ways that Parker and Winsberg have brought out. Both methods in the uses I am focused on employ a stand-in, a study system whose results are to be generalized to a target system. In both cases the justification for that generalization goes by way of establishing relevant similarity between the study and target systems, of whatever sort, by whatever means. Both experiments and simulations are run. That is, they are dynamical processes initiated by the functional equivalent of an ON switch. In both experiment and computer simulation these processes are concrete. In experiment this is obvious; for example, Rutherford's alpha particles are shot at thin gold foil and follow a trajectory dictated by physical law. In computer simulation, the process is a computation governed by dynamical laws encoded in a program. A computation is a physical process. That is, in perfect analogy to an experiment the computer program constitutes a set of dynamical laws that govern the time evolution of hunks of hardware, typically made of silicon. A program is an abstract entity, but so are the laws of physics. What both sets of laws govern are concrete processes. Both methods are interventions in a broad sense. When the switch is flipped on, an initial state – whether this is flying alpha particles and a sheet of gold, or numerical inputs and their associated silicon – is set free to do its work according to the laws.

Both kinds of studies have outputs at the end of the process that are typically called "data", and in both methods the data must be interpreted in order to have results. To do this, one must verify that the intended intervention (physical process, computation) was actually performed, that the data actually reports the desired quantity, and that control for irrelevant factors was achieved. Debugging a program is epistemically analogous to tinkering with a concrete experimental apparatus to make it intervene or measure as intended. In both methods, the claim that the apparatus or program does what is intended is verified by benchmarking, that is, comparing the results to known endpoint values, and to the results of other studies. Results so certified can be used to justify conclusions about the target system, provided a claim of relevant similarity between study- and target- system is justified.

So far the many similarities between the two methods. The epistemic difference is best developed through an example. Folklore suggests that Rutherford's 1911 paper that explained a variety of scattering results via a nuclear model of the atom was decisive. Though this was not true historically or substantively (Heilbron 1968), the experiments by Hans Geiger and Ernest Marsden (Geiger and Marsden 1909, Geiger 1910) showing back deflection of alpha particles which are remembered as particularly well-handled by Rutherford's nuclear interpretation provide examples of a comprehensible and significant type of experiment. It answered the clear question what the deflection pattern of alpha particles shot at metal foil a few atoms thick was. At the time of Rutherford's nuclear interpretation, physicists knew about electrons, their mass and single negative charge, and that the atom was electrically neutral, and so, that because it contained electrons the atom must also contain positive charge. However, they did not know how the positive charge was distributed. J.J. Thomson's "plum-pudding" model dominated, and in this picture the positive charge was uniformly distributed over the atom.

In the folklore version of the story, the fact that alpha particles were scattered through a wide angle when they were shot at a very thin metal foil, and even sometimes deflected all the way back, could only be explained by supposing the atom had a nucleus, because otherwise nothing in the atom would have enough density or charge to deflect the hefty alpha particle that strongly. In the true version, this experiment, even when combined with all of the other scattering phenomena Rutherford's model could explain, was not taken to be decisive, in part because investigation of the atom via scattering had multiple unknowns concerning the structure of the atom, and the scattering properties of the projectiles. Other aspects of the structure of the atom than the distribution of its matter and charge also affected how it would scatter alpha particles. For example, Thomson's model, which had only compound scattering, could explain the back deflection if the radius of the atom as a whole was exceedingly tiny, and while Rutherford evidently thought this was implausible, atomic radii had not yet been measured.[2]

To see the difference between one of the alpha-scattering experiments and an other-things-equal simulation it is helpful to distinguish knowns and unknowns in the former. Many relevant matters in addition to those listed above were already settled. In addition to knowing the mass and charge of electrons, they knew that alpha particles were helium atoms stripped of their electrons and having a +2 charge, and experimenters had the ability to collimate beams of them to shoot at very high speed at small targets, and to make a foil thin enough that an alpha should be meeting atoms only a few at a time. By the time of Rutherford's interpretation, atoms were known to have a number of electrons that was, conservatively, no more than ten times the atomic weight of the atom, a matter highly relevant to whether electrons sprinkled over the atom would have the heft to deflect an alpha strongly.

The first step to constructing a simulation that is epistemically equal to an alpha-scattering experiment is to take all of the things the experimenters knew and did not know and suppose that the simulationist has the same epistemic status toward those matters. The simulationist knows those things the experimenter knows, so, ideally, she can program her model to fulfill them, to work the same way. For example, she will have simulacra alphas with the right "charge" properties, where that means they respond to simulacra negative charge by changing the analog of their positions in the way that physical charges do, according to the Coulomb force.

[2] Rutherford was proposing a model that included multiple innovations: a nucleus, treatment of the alpha as a point mass, and single collisions rather only compound scattering. Thus, despite his obvious confidence, his argument took the form of an inference to the simplest explanation of a variety of experimental results, rather than a direct argument that his was the only possible explanation. After his 1911 paper was largely ignored he recognized that he would not persuade his colleagues until he derived and tested radioactive, chemical, and spectroscopic predictions of his model. Maybe folklore forgets this period of obscurity for his nuclear hypothesis because it was so short. Rutherford got the project of deriving further predictions started soon after when Niels Bohr joined him in Manchester in 1912 (Heilbron 1968, 300-305).

The key question in constructing this simulation is what we are to do about the experimenter's *unknowns*. Experimenters did not know the radius of the atom, its distribution of positive charge, or number of electrons, among other things. These unknown matters play a role in determining the deflection pattern for alpha particles, and do so without becoming known in the process. When alphas are shot at thin gold foil something happens in the foil and consequently to a semi-circle scintillation screen surrounding the foil, which records the hits and yields the raw data of where the alphas landed, and the experimenters do not need to know the structure of the atom in order for this evolution to occur, or in order to interpret the results as wide-angle deflection of alphas by atoms. In the simulation there has to be an analogous computation or evolution, yielding numbers as data. What should that part of the program look like, and what will the programming decisions be based on? Assumptions will need to be incorporated corresponding to the structure of the atom. Otherwise there will simply be no data about what an "alpha" does in response to an "atom."

No experimenter at the time could be confident about these features of the atom, so the other-things-equal simulationist cannot piggy-back on them. An arbitrary choice about them would make the data resulting from the simulation meaningless. One could program multiple simulations based on a variety of different hypotheses about atomic structure, and this would be a fine thing, but this would be equivalent to what Rutherford and Thomson did in constructing their models and determining what followed from them mathematically. Those types of calculations were not experiments, and could give no answer at all to what alphas actually will do when shot at a thin metal foil.

Whatever an experiment does give us in this case, an other-things-equal simulation on the same question has nothing comparable to offer. If the simulationist were to program something to determine the scattering pattern of "alphas" she would be either begging the question of how alphas scatter when shot at atoms or, at least indirectly, making use of results of previous experiments on that question. For example, previous experiments could give support for a hypothesis about the structure of the atom, which then could be relied on in the programming of subsequent simulations. There could be a simulation that gave a non-question-begging answer to the question, but the programmer would have to help herself to more background knowledge than the experimenter needed, so that simulation would not be epistemically other-things-equal to the experiment.

The kind of unknown that the advantage of experiment over simulation rests on is a key to explaining how experiment can teach us anything at all about the world. These are unknowns that play a role in determining the result, but do not themselves need to be known ahead of time or even described, or even learned via the experiment, in order for us to interpret the results of the experiment as giving an answer to the set question. Geiger and Marsden could blackbox the structure of the atom during and after their experiments, as long as they knew that what was determining the scattering result *was* the structure of the atom, which they could know in part by insuring that the sheet of metal was made of atoms and was only a few atoms thick. The ability to so interpret the results while yet refraining from assumptions about some things in the structure of the atom is indeed what allows the results of the experiment to independently test a variety of models of that structure. The results of an other-things-equal simulation that posited assumptions for the unknowns about the structure of the atom would obviously not provide a test of those assumptions that was independent of them, since the results would be in part determined by those assumptions. Interpreting the results of the corresponding experiment depends on assumptions, of course, the assumptions about what I have called "knowns". But these are assumptions to which the simulationist also helps herself. Holding other things equal, the simulationist must make strictly more assumptions in order to give an answer to the question.

The conditions under which experiments are and are not epistemically superior to simulations, according to this argument, divide neatly along the lines of whether there is or is not a particular kind of unknowns in the question being investigated.[3] We can run this argument for the epistemic superiority of an experiment over the other-things-equal simulation for any case in which there are elements in the experimenter's study system that affect the results and are unknown. In contrast, if there are no such unknowns then there is nothing epistemic to point to that the simulation corresponding to the experiment lacks. If there were nothing unknown to us about the structural properties of the atom or alpha particles that play a role in determining the deflection of alphas, then since we would ipso facto have true beliefs about these matters, the simulationist would have the epistemic means to construct a program based on those beliefs that would yield the same results that the world, as acted on in the experiment, would give.[4] If we know everything about the study system that determines the results of an experiment, and can write a program that perfectly

[3] The Rutherford scattering example invites a reformulation of my argument in the following form, suggested to me by . In a simulation's program there is an input module, an output module, and a module in between that determines what the outputs will be from the inputs. The experimenter's advantage could then be placed in that in-between module, which he gets for free whereas the simulator can only program on the basis of some prior knowledge of how the physical system works, which must be based, at least indirectly, on some prior experiments. While this is a handy way of thinking about my point, and a good heuristic, the division of the programming into the three modules, and locating of the issue in the middle one restricts the point to the special case where the only unknowns are in the middle module. In experiments the inputs, for example, alpha particles, also may have unknown properties relevant to the outcomes, which the experimenter will not have to know ahead of time but the simulator will have to make assumptions about. This is why I prefer to make the contrast my argument depends on one between knowns and unknowns wherever they may occur in the system.

[4] It may seem that we do not need to suppose the simulationist has knowledge or even true beliefs about the study system in order to avoid epistemic inferiority of simulation to experiment, as long as her simulation will yield all the same results as the corresponding experiment would get. Constructing a simulation that mimics the experiment's results without even trying to correctly capture the way the experimental study system works would seem to require either a lot of luck or be completely post hoc to the experiment, but leaving that aside this scenario violates the other things equal condition by changing the question. To see this, consider what is required in order for the experimenter to be giving an answer to the question about what alphas do when shot at atoms. He must do preparation that gives reason to believe the study system has atoms set up a few deep. Otherwise the results of the experiment tell us only what alphas do when shot at X, where X could be a piece of cow dung or asphalt for all we know. So too, a simulation that merely gets to results matching those of the scattering experiment can only give an answer to the question what happens when alphas are shot at a substance that behaves internally in a way similar to the rules of the program.

embodies that knowledge – which is a separate issue – then there is no added epistemic value at all in doing the experiment.

An other-things-equal simulation can be just as good as an experiment, as long as there is nothing in the experimental system that affects the results and is unknown to us. This is interesting in principle, but something we can be about as sure of as we are of death and taxes is that in any physical system there are factors we do not know about that affect the results, however slightly. However good our knowledge of the system is, that knowledge is never perfect or complete. It follows then that the best simulation real human beings can make is never just as good epistemically as the corresponding, other-things-equal, experiment on the same question. This is a comparative conclusion, of course, a comparison between experiment and simulation on a fixed question, given fixed background knowledge. It does not identify any absolute limits on how much or how good is the knowledge simulation, or a given simulation, can achieve, and so does not cast general doubt on the epistemic usefulness of simulations. The practical consequences this principled conclusion does have will be discussed in the last section.

4. ONE STEP BEHIND

One might think that this argument supports the view that the advantage of an experiment is that its study system is made of the same stuff as the target. Geiger and Marsden had the gold itself, where the simulationist has nothing. However, material similarity per se is not essential to the argument, for two reasons.

First, I focused on the question of internal validity, or verification, of constructing the study system (experiment or solver) and knowing that it is doing or measuring what one intends, rather than external validity, the justification of extrapolation of the results to other like systems. And the problem for the simulator is not that she has something that is inferior to matter to work with because it is less similar to the target system, or similar in a different way. The problem is that she has nothing at all to use to decide non-arbitrarily how to complete the construction of a program that will give an answer to the question, even in the study system, unless she relies on some background knowledge that the experimenter does not need to have.

Secondly, though material similarity does give the experimenter a leg up in the ability to construct a stand-in that is non-arbitrary, this advantage comes from the fact, well-described by Guala, that material similarity allows us to infer *indirectly* that the stand-in is answering the question of interest, without knowing how it does this. Matter is not the relevant matter, because such indirect support can be gotten by other means, and it is gotten regularly in simulations. In a simulation, a solver is constructed using background knowledge, a model whose broad assumptions we know and are justified in believing have relevant similarities to the target system, but whose consequences for the question of interest, and what exactly determines those consequences, are not known. This is exactly analogous to knowing that a sample of gold contains what will determine the deflection of alpha particles when shot through atoms, without knowing all the structural features of the atom.

Both a simulation and an experiment are run in order to draw out the consequences of intervention on a study system whose similarity but not all of whose features have been established. That is why each method can yield new knowledge. Both kinds of new knowledge are unknown consequences of things we already know or are acquainted with, one logical, the other physical consequences. The discovery of the logical consequences of a set of assumptions and a program is only trivial in the sense Simon's quote above suggests if the discovery of the physical consequences of an intervention on relevant matter is trivial. That would be so only if we thought that experiments were trivial.

It is not because it uses similar matter that experiment is superior to simulation. Under the description just given, a simulation can answer its question with an epistemic force and informativeness equal to that with which an experiment answers its question. However, holding background knowledge constant, a simulation cannot match experimental quality on the same question, except in the case where our background knowledge is perfect and complete. For every given question there can in principle be imagined both an experiment and a simulation that will answer it to equal satisfaction. But the experiment and the simulation that answered that question could not be otherwise equal; the simulation would have to be supplied more information than the experiment would, in our case information about the structure of the atom, in order to get an answer at all. The source of experiment's epistemic superiority to simulation is that simulation is always *one step behind*.

5. PRACTICAL CONSEQUENCES

Supposing I am right that experiment is epistemically superior other things equal, why does it matter when our options in practice actually are constrained pragmatically and our choices are not typically between otherwise equal studies? Obviously, we cannot do a total climate experiment that will tell us what we want to know in time for it to be helpful, and should not detonate nuclear missiles when we have signed a test ban treaty, or deliberately infect human beings with a disease. Why does it matter that these experiments have higher epistemic value if it is not possible or ethical to do them anyway? Granted that we can never justify choosing simulation on grounds that it is just as good as the experiment, since this is never true, why does it matter when often even in cases where the experiment is feasible the simulations we can do are good enough for our purposes?

The choice of simulation cannot rest on an experiment's lack of epistemic advantage but only on that advantage being outweighed by some cost or perceived impossibility. Scrupulousness about this point has psychological and thereby practical consequences. Our only option is indeed to do the best we can within our budget and moral commitments, but budgets and the boundaries of what is possible for us also change in response to our actions. After driving a rental sports car one might decide that owning one is worth economizing in the rest of one's life in order to save up for it. An analogous point holds for funding science; reducing the number of simulation studies funded might make an experiment on that same question affordable. It might not be easy to take this consideration into account in pairwise comparison of the quality and significance of proposed studies, but it could be applied at a more global level of choices among funding strategies.

Not only budgets, but the boundaries of what is possible for us can change, shrinking or expanding in response to our actions. There are cases in which we face the choice of whether to act to make all future experimentation on a question impossible. Since smallpox was eradicated in 1980 it has regularly been proposed that the official remaining stockpiles of the variola virus that causes it be destroyed in order to reduce the risk of future infections. In recent years studies by the World Health Organization and the American Academies of Arts and Sciences have identified specific, medically beneficial research that could not be done if we destroyed the last known stockpiles of this virus, because we do not have sufficient knowledge of the virus to answer them by simulation.[5] But even if we did not know of specific, practically salient questions that only experiments could answer for us, the naïve intuition is correct that there *exist* questions that we can only answer with experiments on the virus itself, and it would always be rational to take that into account in any decision about whether to destroy it.

What is possible for us can not only shrink by our own hand, but also expand in response to our efforts. The boundaries cannot be established a priori. To some it seems that they can but such a stance can turn out to be embarrassing, and as Hume pointed out this seeming can be explained by habit as easily as by the idea that we have found necessary truths:

> Such is the influence of custom, that, it not only covers our natural ignorance, but even conceals itself, and seems not to take place, merely because it is found in the highest degree. (Hume 1999, 110)

There are a lot of experiments it seems we cannot do, but it also once seemed impossible to see the surface of the moon – until Galileo used a telescope – or to measure the speed of light or the density or rotation of the earth. And who would have thought we could do an experiment in which the structural properties of an invisible entity like DNA would make observable differences? Once we have become accustomed to these as actualities it can be easy to forget that they once seemed impossible or even absurd. Science depends on a healthy suspicion toward our current perceptions of the boundaries of the possible.

Perception that the boundaries of the possible are clear has a practical effect on those boundaries through the psychological mechanism of motivation. Belief that something is

---

[5] The World Health Assembly was due to decide whether to recommend destroying the stockpiles in 2011 but, partly due to these studies, and to what the United States believes it can learn for countering bio-terrorism, postponed the decision until 2014. One might think that eradication of the official stockpiles would make the medical and bio-terrorism research unnecessary, but it is not known whether there are unofficial stockpiles. Most governments lobbying in the decision appear to be committed to destruction of the stockpiles but have different preferences about the timetable.

impossible does not imply that it is impossible, but does mean that one has no subjective reason to make the effort, and thereby that a rational person will not actually make the effort. In fact, we may be attracted to believing that a thing is impossible in order to excuse ourselves from making the effort. You can't win if you don't play, as the lottery people tell us, so belief that a thing is impossible prevents us from achieving it, at least with the kind of possibility making in science that does not come by luck.

Surely, it will be objected, one cannot deny that there exist experiments, and many other things, that will never be possible for us. Surely there do, and there is no need to deny it, but a hazard lies in being sure that we know what they are.

REFERENCES

Geiger, Hans. 1910. "The Scattering of the α Particles by Matter." Proc. Roy. Soc. A 83: 492-504.

Geiger, Hans, and Ernest Marsden. 1909. "On a Diffuse Reflection of the α Particles." Proc. R. Soc. Lond. A 82: 495-500.

Gilbert, Nigel, and Klaus Troitzsch. 1999. *Simulation for the Social Scientist*. Philadelphia, PA: Open University Press.

Guala, Francesco. 2002. "Models, Simulations, and Experiments." In, *Model-Based Reasoning: Science, Technology, Values*, eds. Lorenzo Magnani and Nancy Nersessian, 59-74. New York: Kluwer.

———. 2005. *The Methodology of Experimental Economics*. Cambridge: Cambridge University Press.

Harré, Rom. (2003), "The Materiality of Instruments in a Metaphysics for Experiments," In *The Philosophy of Scientific Experimentation*, ed. Hans Radder, 19-38. Pittsburgh: University of Pittsburgh Press.

Heilbron, John L. 1968. "The Scattering of α and β Particles and Rutherford's Atom." Archive for History of Exact Sciences 4(4): 247-307.

Hume, David. 1999. *An Enquiry Concerning Human Understanding*, T. L. Beauchamp ed., Oxford: Oxford University Press.

Humphreys, Paul. 2004. *Extending Ourselves*: *Computational Science, Empiricism, and Scientific Method*. New York: Oxford University Press.

Morgan, Mary. 2002. "Model Experiments and Models in Experiments." In Magnani et al., 41-58.

———. 2003. "Experiments without Material Intervention: Model Experiments, Virtual Experiments and Virtually Experiments." In Radder 2003, 216–235.

———. 2005. “Experiments versus Models: New Phenomena, Inference, and Surprise.” Journal of Economic Methodology 12(2): 317–329.
Morrison, Margaret. 2009. “Models, Measurement and Computer Simulation: the Changing Face of Experimentation.” Philosophical Studies 143(1): 33-57.
Parker, Wendy S. 2009. “Does Matter Really Matter? Computer Simulations, Experiments, and Materiality.” Synthese 169:483–496.
Rutherford, Ernest. 1911. “The Scattering of α and β Particles by Matter and the Structure of the Atom.” Philosophical Magazine 21(125): 669-688.
Simon, Herbert Alexander. 1969. *The Sciences of the Artificial*. Boston: MIT Press.
Winsberg, Eric. 2009. “A Tale of Two Methods.” Synthese 169: 575–592.
———. 2010. *Science in the Age of Computer Simulation*. Chicago: University of Chicago Press.
Wittgenstein, Ludwig. 1969. *On Certainty*, eds. G.E.M. Anscombe and G.H. von Wright. New York: HarperCollins.