Hans Radder (ed.) *The Philosophy of Scientific Experimentation.* Pittsburgh: The University of Pittsburgh Press, 2003.

This volume of essays is the outcome of a workshop that was held at the Vrije Universiteit Amsterdam. For some time analytic philosophers of science have come to pay more explicit attention to the nature of scientific experiments – a topic that for a long time received less attention than one could expect on the basis of what scientists themselves are doing. The present volume provides a nice overview of some of the topics that are currently being discussed under the header of "philosophy of scientific experimentation." As an entry into these issues, the following quotation from *Hans Radder*'s own contribution serves very well:

> [A]n experimenter tries to realize an interaction between an object and some apparatus in such a way that a stable correlation between some feature of the object and some feature of the apparatus will be produced. If the experiment succeeds, two aims have been achieved simultaneously. First, a stable experimental (or object-apparatus) system has been materially realized; second, it has proved possible to obtain some knowledge about relevant features of the object by observing and interpreting correlated features of the apparatus. In addition, since scientific practice does not consist of isolated experiments performed by solitary experimenters, we have to examine the ways in which individual experiments are embedded and used in broader experimental and theoretical contexts. (p. 153)

Experimentation is that part of the scientific enterprise in which scientists engage materially with the world. This engagement is directed towards empirical knowledge. And this knowledge is interpreted, represented, and possibly extended at a theoretical level. Accordingly, one sees arising many levels at which to introduce philosophical enquiries into the nature of scientific experiments. Let us take each of them at a turn, and see what kind of picture can show up from the collected essays.

Materiality.

When experimenting, scientists are forced to leave behind a contemplative stance towards nature. (As is well known, an illustrious forefather of the philosophy of experimentation even spoke about torturing nature.)

Scientists do not get their facts cheap, sitting in a chair, experiencing nature – they hunt for phenomena. Scientists construct traps to catch evidence, they manipulate what they catch, pulling here, pricking there. In short, they confront nature equipped with a host of instruments built to that end. Radder speaks of object-apparatus systems (see quote), *Rom Harré* about apparatus-world complexes. Some of these instruments might function as mere magnifying glasses, but clearly not all of their functions can be captured through such an easy and soothing metaphor. For someone starting from a traditional philosophical perspective, the following question needs to be raised: what kind of nature is left in these "constructed" phenomena?

Peter Kroes characterizes the traditional answer as based on the conviction that the observed phenomena are natural because of the logical separation of *artificial* initial and boundary conditions from the *intrinsic*, natural dynamics of the system involved. By focussing on this distinction, we can introduce some kind of possible classification in the "naturalness" of the phenomena (and concurrently in the kind of instruments used). In some experiments scientists set up a combination of initial and boundary conditions that occurs occasionally without human intervention, in which cases natural behaviour is only *imitated* (possibly to get a better look at it – this is probably the closest one can get to a contemplative ideal of scientific activity). In other experiments, they set up new boundary conditions for naturally occurring phenomena, in which case they are *influencing* natural behaviour – one could say that they are not only looking at, but also playing with the phenomena (possibly to get practically useful results, or to test a theory in some of its more esoteric consequences).[1] Further down the scale of naturalness, scientists are setting up particular combinations of initial and boundary conditions that had not yet occurred on earth (or possibly anywhere in the universe) – one could say they are *constructing* phenomena. But even in this last case the phenomena can be thought to exhibit a natural behaviour, due to their intrinsic dynamics. As Kroes puts it: if we are dealing with some kind of

[1] One could claim that all interactions occurring within measuring instruments form a special subclass of these, in which one wants to ascertain which initial conditions hold in nature, by holding fixed all relevant boundary conditions. E.g. one tries to measure temperature by using a thermometer that is constructed in such a way that its behaviour will depend crucially only on the temperature of the environment.

creation, it is creation in the weak sense of creating instantiations of phenomena, and not in the strong sense of literally creating the properties of the phenomena. However, one should notice that a shift of meaning has occurred in this way of putting the issue: are we talking about creation or construction? That we do not create nature could be readily accepted without thereby necessarily denying that we do construct it. And this construction could be understood in the strong sense. Scientists, on this view, are seen as putting together different elements (possibly including social resources) in the construction of the properties of natural phenomena. There would be no neutral separation of the intrinsic behaviour from the artificial setting of initial and boundary conditions. Nature is what we make of it – without thereby claiming that we make it.

Whatever the view one takes on the general possibility (or desirability) of such a separation between artificial conditions and intrinsic behaviour, quantum mechanics infamously denies it (at least in the Copenhagen interpretation). The behaviour of the phenomena under study is inherently dependent on the setting of the apparatus used to measure and monitor. The apparatus-world complex that is set up in these kinds of situations seems to stand in a truly problematic relation to nature, when all conclusions about nature seem to hold only conditionally (*if* such and such a material setup is created and activated, *then* such and such phenomena will be displayed). Rom Harré, in his contribution, tries to sketch an ontology that would be able to deal with this problem, central to which is the concept of "affordances." But there is a more general lesson to this, on Harré's views. Any classical field theory already requires the abandoning of a traditional substantivist ontology for one in which there is room for non-occurring properties, such as the tendencies of test bodies to accelerate *if* placed in such and such positions; in short, one generally needs ascription of dispositions (to be sure, backed up by causal powers that account for these dispositions). From this perspective, the conditional sentences characterizing the quantum mechanical relation between material apparatus and nature come out to be only a special subclass of this more general scheme. What remains special about them is the human component involved in affordances – what is afforded would not have existed without human action to bring it into being.

As Zen Buddhists probably can testify, one person's contemplation is another one's action. What do we count as a material intervention? Where to draw the border of experimenting "on nature"? *Mary Morgan*'s excellent essay on "Experiments without Material Intervention" broaches these questions - whereas *Evelyn Fox Keller* traces some of their recent history. We cannot deny living not only in a material world, but also in a virtual one. Use of computers in scientific research is becoming very much widespread, in almost all disciplines, and apparently with good reasons; of course there are enormous computational advantages, but there may be more gains to be had: by manipulating artificial computer-systems unambiguous results can be obtained (whereas "real-life" experiments always give more or less dirty results - only vaguely reminiscent of theoretically cherished exact numbers), and moreover all relevant variables are guaranteed to be completely controllable. As Morgan points out, using a study on the strength of bones as an example, there is still an important distinction to be made *within* the class of artificial computer-systems: they can be built to resemble some particular 'real' systems as close as possible in their initial conditions (by using as much empirically gathered information as necessary), or they can be built starting from certain idealizing assumptions (using some selected empirical information). Both kinds of systems can be manipulated on the computer, simulating their "experimental" behaviour under particular manipulations - resulting respectively in what Morgan calls "virtually" and "virtual" experiments - having all the benefits associated with the use of computers. The epistemic status of the respective results will differ however, due to the different status of the input systems with respect to their materiality. As Morgan puts it, the model systems of virtually experiments can be *representative of* similar things in the world (i.e. the semi-material computer bone is a representative one), whereas the model systems of virtual experiments are at most *representations of* another kind of things in the world (i.e. the idealized computer bone is a possible representation of material bones) - in the first case we have possible criteria for judging the similarity and hence the relevance of the obtained results, whereas in the second case we are presented with an inferential gap, excluding the possibility of a direct answer as to why the results should be considered relevant.[2] Of course, in both cases, the dynamics ruling the simulations is pre-programmed, and here lies another important epistemologically relevant factor - arguably *the* most important factor - separating material experiments from non-material ones. Because it is primarily mathematics that governs the evolution of computer systems, all results could have been foreseen. True, this is not always easy - as testified by the need for computers to perform the task - and often one is surprised by the results of mathematical manipulations (why else perform them?). The important point nonetheless remains: the model already contains all resources for explaining the surprising result. Real-life experiments can do something more profound: they can *confound*. An unanticipated result is obtained, and one finds oneself in the impossibility of explaining what happened - for this to be possible, new knowledge is necessary. (This point of Morgan's also shows that the Lord Chancellor's metaphor is quite inapt: I doubt that there has ever been one member of the Spanish Inquisition that was confounded by an unexpected answer, and even less that he would have gone on to revise his own opinions by such a thing; torturing only ends when one hears what was to be heard, and all other answers are neglected.)

Nature's participation in experimentation clearly is what renders science such a powerful enterprise. But the exact nature of this participation remains philosophically contentious, as shown by the discussions concerning constructivism in its different guises.

Empirical Knowledge and Causality

Any arbitrary act by any human being is a material engagement with the world. Surely not all human acts count as experimental. Experimental acts are structured towards a regulative goal, i.e. obtaining stable (object-apparatus) systems, and as a result producing empirical knowledge.

Consider what the natural reaction would be to some confounding experimental situation. The investigated system shows some unexpected behaviour: go and look for the conditions that by being tinkered with produce a different behaviour. If one succeeds, one can infer that these conditions are causally relevant for the behavior. If this investigation is

---

[2] As Morgan makes clear, this need not imply that virtual experiments have no value. Their results, while of limited validity for understanding real systems, might nevertheless be "suggestive."

carried out with enough care, a new stable system can be established.

The preceding remarks of course must be reconstructed with much more care, but their gist is clear: experimental acts are structured on the basis of causal principles. What precise form such causal principles could take is investigated in depth by *Jim Woodward*, introducing the very interesting and promising results gathered in the literature on (computer-aided) causal inference from statistical information (crucially using the so-called causal Markov condition) into the philosophy of experimentation.[3] More specifically he tries to spell out the notion of an intervention, and the role it plays in what I just called the structuring of experimental acts. *Rainer Lange*, while disagreeing with Woodward on the question how precisely to understand causal locutions, follows him in exploring this role. More specifically he focuses on the relations that exist between experimental science and technology, the former being an extension of the latter in his views, the point of separation being exactly the "principle of causality." The difference between an engineer and an experimental scientist, as portrayed by Lange, is that the former is only interested in preventing or suppressing possible disturbances that prevent him to reach a specified goal, whereas the scientist sees these disturbances as potentially interesting in their own right; they are worthy of further investigation and can result in the production of new stable experimental systems. The role played by the principle of causality is that for the scientist any difference in effects *must* be traced back to a difference in causes, rather than trying to suppress the difference by all means. (One can expect that the distinction in practice often will remain vague, and Lange stresses that technology and experimental science always remain closely linked – technology potentially benefiting from new stable experimental systems, and at the same time triggering the production of such systems, by laying bare possible disturbances.)

A very interesting argument by Peter Kroes gives a particular twist to the foregoing considerations. As he remarks, all methodology relies on normative statements regarding the proper *functioning* of the experimental

---

[3] One thing that comes out very clearly from this literature is the profound difference that exists between (statistical) information that is gathered by mere passive observation and information that results from active interventions. (Conditioning on the value of a parameter is fundamentally different from holding the parameter fixed at that value.) Once more, an active stance has to be distinguished carefully from a contemplative one.

setup and the actions performed in assessing the reliability and validity of the outcome of an experiment – and hence in answering the question whether true stability has been attained. Normative evaluations in experimental science are deeply pervaded by functional talk. As a result, teleological notions may be as fundamental to experimental practice as causal ones. If one adds, as Kroes does, the conviction that all attempts to reduce functional language to a strictly causal one are no longer viable, a surprising situation arises. Any experimental situation can be described in two *complementary* ways[4]: straightforward structurally, only referring to causal notions and accordingly barring all methodological judgements, or functionally, stressing the purposeful design of technological artefacts but necessarily leaving out the structural story of how this design is actually embodied and used to bring about particular results. An interesting paradox arises since it becomes hard to justify any particular causal description without the possibility of methodological judgments.[5]
Causal and functional notions are both essential in structuring experimental acts, aiming at (empirical) knowledge. By knowing which specific acts have to be performed to achieve stabilization, one knows that the objects experimented upon enable the establishment of stable systems. But often more specific knowledge is attainable, due to the special character of the structuring of the experimental acts – one can interpret properties of the stable system as reflecting some parts of the causal structure of the world (and hence particular measurement results as being due to some specific property of the object experimented upon). But of course, the question immediately arises whether experimental know-how is enough to make possible such an interpretation, whether it is not necessary to introduce an alien – "theoretical" – element. Moreover, some element has been

---

[4] The term "complementary" is deliberately chosen to bring to mind Niels Bohr's ideas on experimental science. First there is his stress on the impossibility of *describing* nature without a prior *intentional* choice of a particular kind of measurement setup. Next, but related, is the idea that the line separating a studied object and the measurement apparatus used in that study can shift: the apparatus can also be considered to be a part of the object under study, but this then necessarily introduces further apparatuses (with accompanying intentional choices). (Radder also points to this link in his introductory essay.)

[5] It needs to be mentioned that such a stress on non-causal notions might be welcomed by Rainer Lange, who advocates an agency interpretation of causality, in which intentional actions form the basis for our causal notions.

conspicuously missing in our discussion of the production of stable experimental systems up to now: the role theoretical insights play in identifying possible disturbances and ways to control for them – and closely related, in judging that a stable system has truly been realized, i.e. deciding that a system is sufficiently closed from (uncontrolled) disturbances.

Theoretical interpretation

Jim Woodward's analysis of the notion of a causal intervention makes abundantly clear that all methodologically sound causal reasoning requires prior causal knowledge. To get the right (causal) conclusions from the result of material manipulations, one needs to have a correct view on the nature of these manipulations themselves. (One has to assume, e.g., that an intervention only has a local effect, not disrupting other parts of the system.) Moreover the selection of the relevant variables and the way to represent them is not causally innocuous – different choices might result in different conclusions. Causal inferences about the nature of some system can only get started with some representation of the system already in place; hence there are always theoretical interpretations co-responsible for structuring the experimental acts. (In his essay, *Michael Heidelberger* remarks that it is precisely in this causal sense that Hanson's claims about the theory-ladenness of all experiments must be interpreted.) A nice example is provided by Hans Radder when he discusses Newton's experiments on the nature of light. One can only understand the way Newton structured his experimental acts, and accordingly the way he reasoned about them, if one takes into account some of the presuppositions about the nature of light that Newton held to be true – i.e. that the nature of light (and color) is not changed by the passage through a prism. A particularly interesting view on this role of the theoretical context arises when one links this episode directly with Woodward's discussion of interventions: the act by which light is dispersed through a prism counts as a *bona fide* intervention for Newton because he believes that the prism only changes the angle under which the rays of light move, but does not have an impact on the way color is produced by these rays. Someone not accepting Newton's presuppositions could always claim that he had not succeeded in presenting a sufficiently closed system – the prism always introducing further disturbances in the experimental system. The stability of Newton's system might be judged to be only apparent, depending on the theoretical background with which one operates.

In his essay, *Giora Hon* stresses that theory enters at two stages in the interpretation of what is going on in an experiment. The role played by the background conceptual system in making experimental acts meaningful (by describing and interpreting the experimental system – e.g. presuppositions about the relation between prisms and light), is entirely distinct from the one in interpreting the outcome of an experiment (e.g. claiming that the prism experiments prove some facts about the relation between wavelength and color). Since these two roles are logically distinct, no circularity need to be involved in the fact that all experiments always start with theoretical conceptions – as long as it are not these conceptions that are among the conclusions from the experiment.

The two roles mentioned by Hon point toward another aspect about experimental practice that is missing from a perspective that neglects the role of theory: the fact that results can be set apart from the particular experimental procedure from which they resulted. The first theoretical role distinguished by Hon is relevant for interpreting the experimental procedure itself, whereas the second is relevant for the interpretation of the result. This second role enables scientists to detach the results from the specific contexts in which they were obtained. Radder argues that the importance of this possibility is witnessed by the prevalence in scientific practice of *replication*, where scientists try to achieve the same results starting from different experimental processes, as opposed to *reproduction*, where the complete experiment is repeated in all its details. It is not hard to see why replicability should be judged to be such a central goal, as its presence allows the importation of experimental results in new empirical investigations, and hence in the production of new stable experimental systems. As Radder argues, a philosophy of scientific experimentation should make room for these non-local features, and avoid the trap of an excessive localism.

Representation

A last angle under which to look at scientific experiments and the resulting knowledge is by focussing on the ways of representing them.

David Gooding presents an interesting essay on the relation between analogue and digital thinking and representing in the interpretation of scientific experiments. He focuses on the abstraction of natural processes that is inherent in considering them through the mediation of scientific instruments. In this process of *'cognitive narrowing'* a reduced number of quantities are extracted from empirical information, giving us a mathematical workable set of data. If the relevance of these data has to be assessed, however, it is often necessary to re-present them in way that they can be interpreted or manipulated by human reasoners; i.e. through words, symbols and images. The central claim of Gooding's essay, countering some claims on the impact of technological science on our human ways of dealing with the world, is that this *'expansion'* of the reduced - digitalized - information is as fundamental to modern science as is its technological *vista*. Moreover, all (technological) experiments still have to start from some exploratory phase, in which ordinary, experientally based human reasoning remains fundamental. One can only start reducing after the relevant set of empirical information has been selected.

Typically, discussions on representation focus on propositional representations of a body of information. In his essay *Davis Baird* tries to present an epistemology suited to handle the idea of "material knowledge," i.e. the kind of knowledge that is incorporated in material things and our know-how about them. No doubt, attention towards the knowledge borne in our material engagement with the world, not necessarily re-presented at a propositional level, is only to be applauded. What some might have doubts about, however, is the fruitfulness of attempts like Baird's to build his materialist epistemology on an extension of what he calls "traditional epistemological categories," in this way modelling the material knowledge on properties specific to propositional categories. Moreover, these epistemological categories are not eternally given, but crafted towards certain (epistemic) needs - needs that might again be specific to the validation of propositional representations.

Maarten Van Dyck
Universiteit Gent