

**Models and Scientific Representations or:
Who is Afraid of Inconsistency?¹**

Mathias Frisch

University of Maryland, College Park

1. Introduction

It appears to be (still) widely assumed among philosophers of science that a successful scientific theory cannot be inconsistent. My aim in this paper is to outline a framework for thinking about scientific theories and scientific representation and to argue that within this framework inconsistent theories present no special problem. The central tenet of the view I will explore is the claim that representation is an essentially intentional notion, dependent on a context of use. This claim has recently been forcefully defended by Bas van Fraassen (Van Fraassen 2008) and in what follows I will take his account as my main point of departure. I will try to make various aspects of van Fraassen's pragmatic account of representation plausible, but my aim here is not to offer a comprehensive defense of this view. Rather I am interested in the consequences of the account for the question of inconsistency and I will show that on this account inconsistent theories are not especially problematic. My main strategy will be to show how inextricably users of models or representations are involved in scientific representations and to argue that it is the active role of the user that makes successful representations involving inconsistent theories possible.

I will proceed as follows. In the next section I will motivate and partially defend a user-dependent account of scientific representation. In section 3 I will look at the consequences of that account for the question of inconsistency. In the context of this last issue I will revisit an example of scientific modeling that invokes inconsistent laws or basic

¹ Research for this paper was conducted partly while I had a fellowship for experienced researchers by the Alexander-von-Humboldt Foundation at the Ludwig-Maximilians-Universität in Munich. I would like to thank C. Ulises Moulines and Wolfgang Pietsch for helpful discussions. I also want to thank Paul Teller and Peter Vickers for detailed comments on drafts of this paper.

equations—models of particle beams in high-energy synchrotron accelerators, such as the Large Hadron Collider (LHC) at CERN.

2. Scientific Representation

2.1. 'No representation without representer'

What is it for a scientific representation to represent a phenomenon? As van Fraassen convincingly argues, representation is an essentially intentional notion that cannot be defined or be reduced to some other non-intentional notion (2008, 7). To call a thing a representation is to say something about its use. In particular, one cannot define representation independently of any user in terms of likeness or resemblance, for several obvious reasons. As has often been pointed out, resemblance is a symmetric relation but representation is not.² Moreover, while perfect resemblance would be much too strong a requirement for representation, partial resemblance is much too easily to be had: arguably for every two objects there will be some respect in which the two objects resemble each other. Thus, partial resemblance cannot be sufficient for representation, for otherwise we would be forced to the conclusion that everything represents everything else. But partial resemblance as necessary condition would be an empty requirement.

Andreas Bartels (2006) has argued that standard criticisms of resemblance accounts of representation focus on the wrong resemblance relation and are hence too quick in their dismissal of this type of account. Instead of construing representation as postulating an isomorphism—and hence a symmetric mapping—between a representation and its target Bartels maintains that we should think of the representation relation as involving a homomorphism, that is, a many-to-one mapping, which is non-symmetric. Clearly, the existence of a homomorphism cannot be sufficient for representation. But Bartels presents cases that suggest that the existence of a homomorphism might perhaps be necessary for representation. A black-and-white photograph, for example, represents many different shades of color by a single shade of grey. That is, there is a homomorphism from the target to its representation. But it seems that one can equally point to examples in which there is a homomorphism from the representation to its target. For example, a less than perfect

² (Goodman 1976) is the *locus classicus* of criticisms of resemblance accounts.

picture of Ives Klein's *Blue* might represent the strongly monochromatic painting through slight variations in shade. In many contexts it will be understood by the users of the representation that the variations in the shade of blue do not represent variations in the target. Thus, the less than perfect photograph will not misrepresent the painting as being composed of different shades of blue. Note also that arguably there will be *some* homomorphism from the painting to its representation, such as a mapping from variations in the surface structure of the canvas onto the much less structured, flatter sheet of paper onto which the photograph is printed. But even though such a mapping might exist between the representation and its target, it does not feature in the representational role of the former.

The lesson I want to draw from this brief discussion is that we can learn next to nothing about the representational role of an object from examining purely structural relationships between the putative representation and its target objects. How an object represents another object and which features of the object are taken to represent its target and in what ways, are questions that only find answers in a specific context of use. Thus, the "*Hauptsatz*" of van Fraassen's account of representation is that "*There is no representation except in the sense that some things are used, made, or taken, to represent things as thus and so.*" (2008, 23, italics in original)³ In particular, this implies, as van Fraassen himself stresses, that there can be no 'natural representations'—no naturally produced objects or phenomena that represent other phenomena without being used by someone to represent. Independent of its actually being used as a representation a picture has no representational content: "*To call an object a picture at all is to relate it to its use.*" (2008, 25, italics in original)

One might think that computer data or genes present obvious counterexamples to van Fraassen's *Hauptsatz*. An electronically saved file of this paper represents the paper and a piece of the genetic code represents some part or property of an organism without having to be used by anyone to represent the respective targets, or so one might say. But the pieces of computer code have their representational content only derivatively within a context of use, even if a given piece of code is not as a matter of fact used by any particular

³ See also (Frisch 1998). A similar non-reductive account of representation is defended in (Suárez 2004) and (Suárez 2010).

person on a given occasion to represent some other object, and the genetic code does not represent anything at all. Rather, it is part of a mechanism by which organisms transmit properties to their offspring, without thereby representing these properties, just as the gears of my car are part of the mechanism that transmits the engine-power to the wheels, without representing that power. Be that as it may, even if there were counterexamples to van Fraassen's *Hauptsatz* of this kind, this would leave the claim I am interested here unaffected: the claim that all theoretical representations in the sciences are user-dependent.

All this is not to say that resemblance plays no role in representation, but where it does play a role it does so as a function of the representation's use. As van Fraassen points out, sometimes we identify a representation's subject with the help of selective resemblance between representation and target. Which aspects are important in assessing the likeness between representation and target is given by the context in which the representation is used. Moreover, the *success* of certain kinds of representations, including scientific representations, depends on a selective likeness between the representation and its target. And again which aspects of the representation are relevant in judging its likeness to the target and what counts as sufficiently similar for success depends on the context in which the representation is used and can change from context to context. Thus, one and the same scientific model can provide an appropriate representation of an object in some contexts but not in others.

Representation, then, is best thought of not as a two-place relation but rather as a multi-place relation, which includes a place for the user of the representation and for its context, aim, or purpose. If we take aims and purposes to be implicit in the context, we can construe representation as a four-place relation: *a* is a representation of *b*, exactly if there is some context *c* in which a user *u* uses *a* to represent *b*.⁴

2.2. *No perfect model*

Van Fraassen argues that just as partial resemblance can figure in successful representation, *distortion*, or *selective non-resemblance* plays an important role as well. As

⁴ See also (Giere 2006) for a similar spelling out of the representation relation.

an example from physics he discusses the fact that classical physics represents objects as having sharp boundaries. Already the very idea of the true and exact shape and the true and precise boundaries of a macroscopic physical object might strike one as suspect. More troubling, sharp boundaries in a mathematical model often result in discontinuities or singularities, where the physics used to represent a system's behavior sufficiently far away from the boundaries breaks down. Thus, representing objects as having sharp boundaries could not possibly be completely accurate and non-distorting, at least if the resulting models involve singularities. Nevertheless such representations fulfill an important role in our scientific image of the world and may in certain contexts provide us with the only means to construct useful representation of the phenomena.

Sometimes there are techniques to patch over what happens at such boundaries, but the result is often still not a single unified representation of the system. As Mark Wilson argues (Wilson 2008), our representational practices in classical physics frequently have to distort to be successful at all in ways that represent phenomena by partially overlapping yet in some sense incompatible 'theory façades' that break down at the fault lines between the façades. The situation Wilson presents is in certain respects analogous to multi-perspectival paintings by Picasso, that bring together on a single canvas different perspectives on different parts of the human body, without, however, being able to combine these multiple perspectives into yet another unifying perspective (see also van Fraassen 2008, 38).

However, I want to focus here on a different cluster of reasons for why distortions play a central role in scientific representations. Paul Teller (Teller 2001) has argued forcefully against what he calls "the perfect model model"—that is, the view that our best scientific theories provide us with complete and perfectly accurate models of physical phenomena, or at least that physics is progressing toward and aiming at developing ever more complete and accurate and non-distorting models of the world.⁵ Teller argues that this view of physics is mistaken and emphasizes in its stead the importance of highly idealized models. Idealized models distort in that they represent only some aspects of the physical system modeled while leaving out other aspects and may purposefully

⁵ See also (Teller 2004).

misrepresent some of the aspects represented. According to some of Teller's arguments, which I want to amplify here, distorting models at the very least play an important explanatory role and would not be rendered explanatorily superfluous by complete and perfect models. But there are even stronger arguments that suggest that the idea that physics even in principle presents us with perfect models of the phenomena is a myth and that all scientific representations are distorting.

Teller illustrates the importance of distorting models by pointing to different ways in which physicists model different aspects of the behavior of water. Continuum models can correctly represent the wave behavior of water, whereas particle models are used to represent diffusive behavior. While both kinds of model represent certain aspects of the behavior of water sufficiently accurately, neither type constitutes a perfect model of water that can successfully represent all of its properties. And both types of model manage to capture aspects of the behavior of water by distorting, by representing water either as a continuum or as a collection of classical particles. The example, thus, supports the claim that given the way our world is, for many physical systems there is no single model of the system that allows us successfully to represent the kind of system in all circumstances and we need to employ different, and even in some sense incompatible models to represent successfully different aspects of the system's behavior.

Now, one might reply to Teller's example by arguing that in addition to these two types of models there exists a third type—quantum mechanical models—that do provide us with perfect, non-distorting representations of water and that can unify the two classical types of model by showing how both can be approximately derived from the correct and complete micro-theory by taking appropriate limits. This reply can be read as arguing either against the claim that the idealized models are *explanatorily ineliminable* or against the claim that *all models* in physics are distorting. I want to focus on the issue of explanatoriness first.

Let us provisionally grant for the moment that present-day quantum mechanics provides us in principle with an immensely complex yet accurate—indeed, perfect—model of water. Nevertheless, as Teller argues convincingly, even if we imagined that we were given the solution to the Schrödinger Equation for a system of 10^{25} variables, this would not provide us with an explanation of macroscopic waves in a body of water, since being

able to grasp what this solution says about the behavior of the body of water goes far beyond what is cognitively possible for us humans. Explanation and understanding are inherently pragmatic notions—explanation is always *explanation for us*—and we absolutely require the more idealized classical models to be able to understand the water’s behavior.

I made what is essentially the same point in (Frisch 1998) by contrasting classical Newtonian explanations of the collection of balls known as “Newton’s cradle” with a putative quantum mechanical account of the system of five balls. I also argued there that there are additional reasons not to take a putative quantum mechanical account as providing a satisfactory explanation. Scientific explanations answer “what-if-things-had-been-different-questions” (see Woodward 2003); that is, a scientific explanation embeds the explanandum phenomenon into a pattern of counterfactual dependencies. In the physical sciences explanations provide mathematical models of a phenomenon that embed it into a pattern of functional dependencies, which informs us how features of the explanandum vary with the values of the variables used to represent the phenomenon. In a sense, this account puts rather weak constraints on what counts as an explanation: functions are easy to come by. What, on my view, distinguishes putative explanations from one another is not so much the question whether a given account is an explanation or not—even Moliere’s dormative virtue is weakly explanatory, I want to claim—but rather the relative strength of an explanation. The goodness of an explanation depends on its accuracy, strength, and simplicity. An explanation is stronger, if it includes more factors that make a difference to the occurrence of the explanandum. The consideration of simplicity pulls in the opposite direction: An explanation is simpler if it does not add irrelevant details. Both these conditions, as well as what counts as sufficient accuracy, are context-dependent. Including certain variables in a model of a phenomenon may be explanatorily important in one context, while the same variables may be adding irrelevant details in another. The best explanation is one that scores best on some weighted average of these criteria, where again there is no context-independent algorithm for computing this average.

A consequence of this account is that an explanation that is more accurate than all of its rivals need not be the best explanation of a phenomenon. In fact, in most contexts it will be the case that a microscopic quantum-mechanical account of the state of all the molecules

composing the water will be taken to offer much too many unnecessary details for successfully explaining the water's wave-like behavior. Thus, in most explanatory contexts a full microscopic quantum mechanical model would be explanatorily inferior to the classical model *even if per impossibile* we were able to cognitively grasp the details of the former model. What is crucial for an understanding of the behavior of water waves (in most contexts) is an understanding of the general patterns that transcend the details of the given case—patterns that arguably would be lost in the minute details of a quantum mechanical model with in the order of 10^{25} variables and that transcend the details of a putatively quantum mechanical micro-model. Thus, even if we were to grant that quantum mechanics provides us, at least in principle, with a perfect model of the behavior of water, this would not render the distorting models superfluous: in many contexts the distorting models still provide us with the best explanation of the behavior of water and their explanatory success depends precisely on the fact that these models are highly idealized and distorting.

But in fact we have granted much too much to the objection, for we do not actually possess a perfectly accurate and complete quantum mechanical model of wave or diffusion phenomena. This brings us to the second issue—the foundationalist assumption that there exist perfectly accurate, fundamental models, constructed with the help of our most fundamental theories, underlying the idealized higher-level models we use in practice. I now want to challenge this assumption. First, quantum mechanics has its own limits in scope and accuracy—current quantum mechanics, too, distorts and is not the final and correct theory of the world. Second and more importantly, even if we were to believe that present day quantum theory was exactly correct wherever it can be applied, it still would not provide us with models of the macro-phenomena at issue here. We have so far imagined that we were somehow given a quantum-mechanical model of macroscopic bodies of water, but that is of course only an impossible fiction. To actually construct any such model, we would have to solve the Schrödinger equation for on the order of 10^{25} variables—something that is impossible to do and far beyond our computational capacities.

At this point one common reply is to insist that even if it is impossible actually to solve the Schrödinger equation for macroscopic systems, the theory nevertheless provides us with models of arbitrary complexity. The equation defines a class of models, many of

which we of course never construct explicitly. Indeed, any physical theory has many, many more models than the ones scientists have actually constructed and actually used.

Quantum mechanics contains a model of the hydrogen atom, of Bose-Einstein condensates, but also, one might argue, of any arbitrary system. If a solution of the Schrödinger equation exists for certain arbitrarily complex initial conditions for systems consisting of 10^{25} particles, then simply in virtue of being in possession of the equation we thereby also are given a model of such systems, even if we do not know how to—or are practically unable to—explicitly construct this solution.

But the mere fact that an equation has solutions in addition to the ones actually constructed does not imply that whenever we possess a theory we thereby also are in possession of a large range of models of arbitrarily complex systems governed by that theory. That quantum mechanics provides us with a model even of macroscopic bodies of water is supposed to follow from the claim that among the set of solutions to the Schrödinger equation are ones that are structurally similar to bodies of water. Yet, as we have learnt from van Fraassen's account, no structural relationship between a model and a phenomenon can on its own suffice to make the model a representation of that phenomena (see also van Fraassen 2008, 250). Rather something is a representation only if it is *used* to represent a thing. But since we do not even actually *have* the quantum mechanical initial state of the system of water let alone a solution of the Schrödinger Equation for that system—since there is no way for us to pick out the appropriate solutions from the class of solutions defined by the Schrödinger equation—we cannot *use* the putatively existing solution to represent anything. Thus, as van Fraassen warns us, we have to be careful about an “illegitimate slippage from ‘there exists’ to ‘we have’” (van Fraassen 2008, 233). While there may exist solutions to the Schrödinger equation for systems of 10^{25} variables, simply writing down the general form of the equation does not imply that we thereby have all of its solutions of arbitrary complexity.

But this conclusion seems highly counterintuitive. In accepting Newtonian physics, say, are we not committed to the claim that the theory successfully applies to planetary systems yet to be discovered and systems of billiard balls never explicitly modeled? Any theory's domain, it seems, extends well beyond the class of phenomena for which we have actually constructed models. To answer this puzzle, it seems to me that we need to

distinguish carefully between two ways in which we might think that a theory's representational reach extends beyond the systems actually modeled.

Van Fraassen's own example of a theory representing phenomena that we have not modeled is that of a colony of bacteria located somewhere in Antarctica long before the first humans appeared on Earth. He asks whether we can say that a theory of exponential growth adequately represents the evolution of this colony, even though by hypothesis no model for this particular phenomenon was ever offered. His reply is that the theory is adequate if among the solutions to its equations is one defining a structure that would satisfy the relevant constraints on adequacy if it were used to represent the colony's evolution. The worry, of course raised by his own context- and user-dependent account of representation is whether this counterfactual has reasonably well-defined truth conditions. It seems to me that the answer is 'yes' in the present case, since scientists actually and as a matter of fact use models of bacterial colonies to represent their growth and arguably this practice sufficiently constrains how we would represent the Antarctic colony were we to do so. That is, scientists actually depict bacterial colonies through what van Fraassen calls "data models" that are appropriate for a representation of the colonies' evolution in terms of exponential growth models; and scientists actually use the latter models to represent bacterial colonies. Arguably, this practice sufficiently constrains what it *would* be to provide a data model of the Antarctic colony—that is, what it would be for us to selectively structure the phenomenon in a way that is relevant to exponential growth theory. As van Fraassen correctly emphasizes, however, the notion of relevance here is relative to a user and a specific context of use:

*There is nothing in an abstract structure **itself** that can determine that it is the **relevant** data model, to be matched by the theory.* A particular data model is relevant because it was constructed on the basis of results gathered in a certain way, selected by specific criteria of relevance, on certain occasions, in a practical experimental or observational setting, designed for that purpose.
(2008, 253, italics and emphases in original)

There is a well-defined answer to what it would be to depict the Antarctic colony and embed its data model into a model of exponential growth only because the situation so closely resembles phenomena we have actually modeled and which can therefore provide

appropriate criteria of relevance. The situation is dramatically different in the case of the putative micro-model of water—for the solution to the Schrödinger equation for 10^{25} variables which we have posited exists, but which do not actually possess and cannot as a matter of fact use to represent macroscopic bodies of water.

In this and the previous section I have discussed various ways in which scientific representation is inextricably user- and context-dependent. It is this user-dependence, which allows for inconsistent sets of assumptions to be used in providing us with successful representation of the phenomena within a domain, as I now want to argue.

3. Inconsistency

Traditionally, many philosophers of science have taken it to be a necessary condition that scientific theories have to be consistent and even those philosophers who do allow for inconsistent theories to play some useful role in science take them to be at best preliminary stepping stones in the development of consistent theories. There are three worries one might have about inconsistent theories. First, inconsistent theories could not possibly be true.⁶ Thus, if the aim of science was ultimately to arrive at true theories about the world, an inconsistent theory could never be the end-product of a scientific development. Second, and even more troubling, inconsistent theories raise the threat of logical anarchy, since every sentence can be logically derived from an inconsistent set of sentences. That is, inconsistent theories do not seem to exclude anything. A third worry is that an inconsistent theory is uninterpretable.⁷ Suppose we understand interpreting a theory as saying ways in which the world could possibly be so as to make the theory true. And suppose we understand saying how the world could possibly be as characterizing a possible world, even if understood only metaphorically, in which the theory is true. Then it would seem to follow that an inconsistent theory is uninterpretable.

On the account of scientific representation outlined above, however, none of these worries arises. Inconsistent theories cannot be true (if we do not allow for the world itself to be inconsistent), but neither are the consistent theories we possess true. Once we are no longer in the grips of the ‘perfect model model’ and realize that all our theories provide us

⁶ Nor could they be true about what is observable.

⁷ This worry was suggested to me by Paul Teller.

with only partial and partially distorting representations of the phenomena in their domain, the idea that the aim of science is ultimately to arrive at true theories becomes untenable. To be sure, a consistent theory might still possibly be true, in some sense of 'possible' while an inconsistent theory cannot, but this does not constitute a difference that is important to our representational practices. If often (or even always) our best models of a phenomenon are partially distorting representations then it does not matter whether the theory with the help of which we constructed the representations provides us with a set of satisfiable sentences or not—as long, that is, that inconsistent theories do not result in logical mayhem. Put differently, if the end-product of a scientific development can be false—that is, if the best explanation of a phenomenon in a certain domain can be one that is provided by a theory that is strictly speaking false—then the necessary falsehood of an inconsistent theory presents no hurdle to such a theory providing us with what may be the best explanation of a certain phenomenon.

The second worry is that inconsistent theories result in logical mayhem. An inconsistent theory has no models, in the model-theoretic sense, and there seem to be no well-defined procedures for constructing representations of the phenomena with the help of inconsistent theories. But once we give up what Nancy Cartwright has called the 'vending machine view' of scientific theories, this worry, too, does not arise: while there are no universal, purely logical rules that govern how representations are constructed with the help of an inconsistent theory, there can be context-dependent and non-arbitrary additional constraints that users of a theory employ in constructing models.

Now, strictly speaking this reply relies on an extension of the account of representation presented above. So far we have stressed the role of the user both in depicting phenomena in a structured way and in using models to represent a phenomenon—that is, in how we represent the world and in forging model-world relations but not in the construction of theoretical representations. But given the central role that users of representations occupy in the account, it is a natural extension also to take users to play an important role in the construction of models with the help of the theoretical resources available. We have already seen that the vending machine view fails, since the set of structures defined by a set of basic equations cannot be identified with the class of

representations provided by a theory. The additional claim now is that users also play a crucial role in the construction of representations.

The answer to the second worry, then, is that there are additional context-dependent constraints that users of a theory employ in constructing representations with the help of a set of basic equations, and that the existence of such constraints allows even for formally inconsistent sets of equations to play a useful role in science.

What about the third worry? We can grant, at least for present purposes, that there is something right about its core intuition that we understand the content of a theory by exploring possible worlds or possible systems in which the theory's basic equations hold. This is in fact how physics students come to learn a physical theory: by exploring paradigmatic or exemplary systems to which the theory's laws apply. But this process of understanding and interpreting a theory need not happen in one fell swoop, as it were. Rather it can begin with very simple systems which satisfy only a subset of the theory's equations—a subset which can be consistent even in the case of inconsistent theories—and then proceed to more complex systems, which with the help of context-dependent constraints integrate different, possibly formally mutually inconsistent parts of the theory. As we will see in a bit more detail below, as long as there are context-dependent constraints on allowable inferences we can even explore what are in some sense impossible worlds and ask how these worlds evolve or how the world would change if the values of certain quantities would change.

At this point it will be useful to have an actual example to consider. Thus, I briefly want to review what I have argued elsewhere is a case of modeling a phenomenon with the help of an inconsistent set of assumptions: models of particle beams in synchrotron accelerators, such as the Large Hadron Collider (LHC) at CERN (see, e.g., Steinhagen 2007, 23-25).

In the LHC a proton beam is accelerated by external electric fields and is held on a circular orbit with the help of magnetic fields. The equations used in modeling the beam are the Maxwell equations and a Newtonian equation of motion with the force acting on the charge being given by the Lorentz force due to the *external* electromagnetic fields. Alternatively, one might write down an equation in which the momentum-change of the charge depends on the *total* field, including the self-field associated with the charge, but

there seems to be no fully satisfactory way of including self-interaction effects into in a fully consistent and conceptually unproblematic classical theory (see Frisch 2005). In any event, the equation of motion used in modelling synchrotron radiation ignores the self-fields. The Maxwell equations are used to determine the electromagnetic fields due to the bending magnets and the accelerator components with which the protons interact. The Lorentz force law is used to determine the motion of the proton beam via the relativistic analogue of Newton's equation of motion, $dp^{\mu}/d\tau = F_{ext}^{\mu\nu}v_{\nu}$. Furthermore, one assumes energy-momentum conservation. These assumptions together, however, yield an inconsistency.⁸

Does this mean that the theory of classical electrodynamics is inconsistent? I think it does not matter for our purposes here how we answer this question. What I am arguing here is that the mathematical assumptions used to model a wide class of phenomena in classical electrodynamics form an inconsistent set. The inconsistency would lead to problems for a non-pragmatic account of scientific representation, independently of whether the set of assumptions are taken to constitute the core of a 'theory' or not.

I want to distinguish between two types of inconsistency: First, the set of equations invoked in constructing a model may be formally mutually inconsistent; and, second, while the equations themselves are consistent, they are inconsistent with the existence of the very kind of physical object to which they are applied. The inconsistency in the case of synchrotron radiation is of the second type. The Maxwell equations and the Lorentz force equation have solutions that satisfy energy conservation—solutions for continuous charge distributions or charged dusts. But they do not have any solutions for discrete finite charges—that is, solutions where the motion of a discrete charge is given by the Newtonian Lorentz force equation of motion—that satisfy energy-momentum conservation. The problem is that an accelerating charged particle should, according to the Maxwell equations, radiate energy. Yet this energy loss is not accounted for in the energy-balance implied by the Lorentz force equation of motion, which only takes into account *external* fields but not the charge's own field—the so-called 'self-field.' Thus, the Lorentz force

⁸ I have discussed this inconsistency in (Frisch 2004) and (Frisch 2005a). My account has been criticized by (Belot 2007) and (Muller 2007), and criticized as well as partially defended by (Vickers 2008). I respond to some of the criticisms in (Frisch 2008) and (Frisch 2009).

equation, applied to a discrete charged particle, and the Maxwell equations are inconsistent with energy-momentum conservation.

One common view is that when physicists employ what appears to be an inconsistent set of equations in constructing a model, closer analysis will usually reveal that what is in fact used is a restriction of the equations at issue to a fully consistent subset of equations. This view does not, however, fit well how the beams in synchrotrons are modeled. While there is a sense in which any model of a particular experiment is consistent, models are used within the context of an inconsistent set of background assumptions that govern the model's application. To see what I have in mind here, consider as an analogy one of M. C. Escher's famous lithographs of impossible situations, *Waterfall*, which depicts a waterfall powering a waterwheel. At the bottom of the waterfall the water flows apparently downhill in an aqueduct, which, however, ends at the top of the waterfall. The picture considered on its own is obviously not impossible—it consists of a series of perfectly possible (because actual) lines on a sheet of paper. But given the context and the representational conventions within which it quite consciously is meant to be viewed, it suggests an inconsistent or impossible state of affairs: the picture is a representation of a fictional stream of water which is represented both as continuously flowing downhill and as flowing in a closed loop. Yet this impossibility does not imply that we cannot meaningfully explore different aspects of the situation depicted by Escher. Rather, the background assumptions governing our representational practices allow us to ask well-defined 'what-if' questions. For example, we can ask what would happen to a little toy boat placed in the aqueduct (it would continuously go around with the water in the closed loop); or we can conclude with Escher himself that in order to keep the water wheel as perpetual motion machine running indefinitely, one would periodically have to add water to compensate for evaporation. Given the representational context, the lithograph represents a water cycle as having inconsistent properties, but this does not lead to logical mayhem.

I want to submit that Escher's lithograph is analogous to the models of the particle beam in synchrotron accelerators in important ways. Here, too, the model itself is not impossible. The model consists of a particular external field, a specific particle orbit, and the protons' radiation field, which on their own are no more incompatible with each other

than the lines on Escher's drawing. But, as in the case of Escher's *Waterfall* the model is constructed in a context and with the help of principles that imply that the model represents a physically impossible state of affairs, namely protons that orbit with constant radius even as they radiate off energy. Nevertheless, as in the case of *Waterfall*, we can ask informative 'what-if-things-had-been-different' questions about the model, such as how the orbit would change if the external field changed or if we had injected Helium nuclei instead of protons into the accelerator. Importantly, all three of the assumptions incompatible with the existence of discrete charges are crucial to the uses to which the model is put: the Maxwell equations allow us to calculate variations in the external field and the protons' radiation field, the Lorentz force equation of motion enables us to probe the protons' orbit, while the principle of energy conservation is crucial to analyzing the decay products that result from collisions between anti-cyclical proton beams.

There are, however, also important disanalogies between Escher's lithograph and the synchrotron model: while the former is intended to represent a fictional situation *as impossible*, the synchrotron model is intended to represent *possible* orbits of charged particles. The lithograph focuses on and draws our attention to the impossibility Escher aims to represent. By contrast, even though the synchrotron model attributes incompatible properties to the proton beam it does not represent the proton beam *as impossible*. Rather as an unavoidable consequence of the mathematical machinery employed, the impossibility is part of the representational background. The situation is as if in the artistic case representational conventions or media required that water in a closed loop aqueduct be represented as always flowing downward, even in the case of a picture intended to represent a still body of water as still.

A second disanalogy is the fact that the synchrotron model, unlike the lithograph, centrally involves magnitudes, which allow us to quantify the departure from consistency, as it were. If one calculates how much energy the particles lose due to radiation in a synchrotron accelerator, one can determine by how much the orbit calculated with the help of the Lorentz force law would have to be adjusted for energy to be approximately conserved. In the case of highly relativistic electrons, the energy loss is non-negligible, while in the case of the LHC proton beam it is negligible. That is, in the case of electrons the Lorentz force equation of motion leads to predictive inaccuracies too large to be acceptable

in many contexts and which have to be remedied through stepwise corrections, while in the case of the much more massive proton beam the initial model represents the beam sufficiently accurately.

(Vickers 2008) describes my account of modeling synchrotron radiation as involving a commitment to the claim that the self-force on a charge due to its own field is zero. But this is not quite right. It is only right as far as the initial calculation of the orbit of the charge is concerned based on the external Lorentz force equation. In the second step, however, the radiation field of the charge is calculated, which implies a non-zero self-force, and it is then checked if the error made in ignoring the self-force in the first step was too large. Vickers also suggests that perhaps the model is simply silent on whether the self-force is zero or not. But without making a commitment on what the total force acting on the charge is, we do not have a well-defined equation of motion from which to calculate the particle orbit. Thus, the model is committed to the claim that the self-force is non-zero in the context of calculating the radiation field but is committed to the claim that the self-force is zero in the context of calculating the particle orbit.

Now, the way I have described the process very much looks like a standard case of approximation: we use an approximate equation of motion (and hence, one that may be formally inconsistent with the remainder of a theory's equations), because it is simpler and introduces only a negligible error. What is interesting about the present case, however, is that there is no acceptable fully consistent set of equations to which the Lorentz force equation could be thought to be an approximation. The obvious solution would be to replace the external Lorentz force equation with an equation of motion that does take the self-field into account. However, as I have argued in (Frisch 2005, 2008), there is no fully satisfactory way of doing this. Once we take charged particles to be subject to the total Lorentz force, including the self-force, questions concerning the structure of charged particles become important. The simplest option is to treat charges as point particles. Yet taking the self-field into account then leads to infinities and these infinities cannot satisfactorily be 'swept under the rug' through a process of renormalization. Models of extended charged particles avoid the infinities, and there are even existence and uniqueness proofs of solutions for the Maxwell-Lorentz equations for certain non-relativistic charge models. Yet there is no known fully relativistic model of extended

charged particles. There are various proposals for taking self-fields approximately into account, chiefly among them in a ‘regularized’ equation of motion, but this equation, too, does not satisfy the logicians’ demand of strict consistency with the Maxwell equations and the principle of energy-conservation. Thus, the modeling practices employing the external Lorentz force equation of motion are not backed up by a fully consistent set of equations to which the Lorentz force equation could be understood as an approximation.

The problem appears to lie with the unhappy marriage of particle and field theories that the contemporary—that is, the post-Lorentzian—microscopic theory of classical electrodynamics attempts to achieve. Extended particles do not fit well with a relativistic field theory, which would require extended objects to have an infinite number of degrees of freedom, but point particles result in singularities. Nevertheless the particle-field framework has proved to be extremely fruitful for modeling electromagnetic phenomena. Putting it another way and adopting Vaihinger’s characterization of the atom, a discrete charged particle appears to be “a contradictory representational knot, which nevertheless is necessary for the calculation of reality.” (Vaihinger , 102)

Conclusion

I have argued that modeling the interaction of classical charged particles with electromagnetic fields relies on a set of inconsistent assumptions. Without emphasizing the role of the user and context-dependent constraints in scientific representations, it is difficult to see how these assumptions could result in scientifically successful predictions. Once we put the user into the picture, however, there is no longer a special problem associated with inconsistent theories or modeling assumptions. If, as I suggested above, we make explicit the role of the user not only in the application but also in the construction of a model or representation, then inconsistent modeling assumptions no longer pose an insurmountable obstacle to our representational practices.

Physicists use the external Lorentz force equation to model the the motion of electric charges, such as the LHC proton beam, in an external electromagnetic field, but they do not use the equation to conclude that the energy change of the charge equals the external work done on the charge, which would imply that accelerated charges do not radiate. That is, even within the construction of a single representation, there can be

constraints as to which inferences are licensed by a given principle and which are not. Put another way, the contexts for the use of a given assumption are individuated rather finely: an assumption can be used to construct one aspect of the representation of a particular phenomenon, without one being committed to whatever implications the assumption might have for another aspect of the very same representation.

References

- Bartels, Andreas. 2006. Defending the Structural Concept of Representation. *Theoria* 55: 7-19.
- Belot, Gordon. 2007. Is Classical Electrodynamics an Inconsistent Theory? *Canadian Journal of Philosophy* 37, no. 2 (June): 263-282.
- Frisch, Mathias. 1998. Theories, Models, and Explanation. Thesis (Ph.D. in Philosophy) -- University of California, Berkeley, May 1998.
- . 2004. Inconsistency in Classical Electrodynamics. *Philosophy of Science* 71, no. 4: 525-549. doi:10.1086/423627.
- . 2005. *Inconsistency, Asymmetry, and Non-Locality: A Philosophical Investigation of Classical Electrodynamics*. Oxford: Oxford University Press.
- . 2008. Conceptual problems in classical electrodynamics. *Philosophy of Science* 75, no. 1.
- . 2009. Philosophical issues in electromagnetism. *Philosophy Compass* 4, no. 1: 255-270.
- Giere, Ronald N. 2006. *Scientific Perspectivism*. Chicago: University of Chicago Press.
- Goodman, Nelson. 1976. *Languages of Art: An Approach to a Theory of Symbols*. 2nd ed. Indianapolis: Hackett.
- Muller, F. A. 2007. Inconsistency in Classical Electrodynamics? *Philosophy of Science* 74, no. 2 (April): 253-277. doi:10.1086/520942.
- Steinhagen, R. J. 2007. LHC Beam Stability and Feedback Control. CERN-AB- 2007-049 B1. cdsweb.cern.ch/record/1054826/files/ab-2007-049.pdf.
- Suárez, Mauricio. 2004. An Inferential Conception of Scientific Representation. *Philosophy of Science* 71, no. 5 (December 1): 767-779. doi:10.1086/421415.
- . 2010. Scientific Representation. *Philosophy Compass* 5, no. 1: 91-101.

doi:10.1111/j.1747-9991.2009.00261.x.

Teller, Paul. 2001. Twilight of the Perfect Model Model. *Erkenntnis (1975-)* 55, no. 3: 393-415.

———. 2004. How We Dapple the World. *Philosophy of Science* 71, no. 4 (October 1): 425-447. doi:10.1086/423625.

Van Fraassen, Bas C. 2008. *Scientific Representation: Paradoxes of Perspective*. Oxford: Clarendon Press.

Vickers, Peter. 2008. Frisch, Muller, and Belot on an inconsistency in classical electrodynamics. *British Journal for the Philosophy of Science* 59, no. 4.

Wilson, Mark. 2008. *Wandering Significance: An Essay on Conceptual Behaviour*. New York: Oxford University Press.

Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. New York: Oxford University Press.