

Renormalization and the Formulation of Scientific Realism

James D. Fraser*

Forthcoming in Philosophy of Science

Abstract

Providing a precise statement of their position has long been a central challenge facing the scientific realist. This paper draws some morals about how realism ought to be formulated from the renormalization group framework in high energy physics.

1 Introduction

Many philosophers of science subscribe to scientific realism. Unfortunately, there is much less agreement about what this doctrine amounts to. My suggestion in this paper is that the renormalization group framework in high energy physics provides a useful case study when it comes to the question of how realism ought to be formulated. The claim that the renormalization group has important implications for a realist view of quantum field theory (QFT) has been mooted in the past by David Wallace (2006, 2011), and more recently by Porter Williams (2017). I develop this line of thought in a broader context here and argue that there are lessons to be learned for the broader realism debate. On the one hand, the role the renormalization group plays in identifying aspects of QFT models we should take representationally seriously supports a local approach to articulating the realist thesis; rather than attempting to explicate how theories latch onto the world in general terms it shows that resources found in particular scientific contexts can be a crucial part of this story. On the other hand, it points to a strategy for responding to Kyle Stanford's "trust argument". Stanford challenges the realist to tell us what of our current theories will survive future scientific progress. While this can seem to be an impossible task in the abstract I will suggest that it may become more tractable at the local, theory-by-theory, level.

The plan for the paper is as follows: section 2 provides an opinionated overview of some key issues surrounding the formulation of realism, section 3 introduces the

*Email: jamesf09@hotmail.co.uk.

renormalization group and explains how it helps substantiate a realist analysis of QFT models, and section 4 draws out some broader morals for the formulation problem.

2 The Formulation Problem

What is scientific realism? A naive answer is that it is the claim that our best confirmed scientific theories are true. It quickly becomes obvious that construing realism in this way renders it a wildly implausible doctrine however. One much-discussed reason is the pattern of theory change found in the historical record. Examples of theories which made accurate predictions in their day but later turned out to be false are legion in the history of science, and while philosophers with realist sympathies have pushed against the idea that this undermines the connection between empirical success and truth entirely, they typically admit that it gives us grounds to doubt that current theories get everything exactly right. There are also ample indications within contemporary science itself that our theories are not entirely veridical. To take a particularly stark example, our most fundamental physical theories, QFT and general relativity, furnish mutually inconsistent accounts of what the world is like, and powerful theoretical arguments suggest that a completely new framework is needed to fully describe gravitational phenomena.

What should the would-be realist replace this naive formulation with? What seems to be needed is a more modest epistemic attitude towards predictively successful theories: something stronger than merely taking them to be empirically adequate, as the constructive empiricist advises, but weaker than believing everything they say about the world. A common move when outlining the realist position in its broad brush strokes is to replace the word ‘true’ in the naive formulation with ‘approximately true’. Ultimately though, this only postpones the problem, as what it means for a theory to be approximately true stands in dire need of clarification. While there is a great deal of disagreement over how a more precise formulation of realism should be developed in detail, some consensus has emerged about the general form it should take. It is widely agreed, amongst both contemporary realists and anti-realist critics, that any formulation worthy of serious consideration must be ‘selective’, recommending we commit ourselves to some parts of a successful theory’s content, while rejecting, or remaining agnostic about, others. The challenge to spell out the sense in which successful theories are approximately true can then be answered, at least in part, by pointing to a subset of their claims about the unobservable which hit their mark.¹ This selective approach owes its popularity to the dominant realist approach to the challenge of historical theory

¹To be sure, many questions remain about approximate truth and its role in the formulation of realism. Some selective realists ascribe the property of approximate truth to individual theory constituents, for instance, so there is a residual puzzle about how to make sense of this notion at this level of particular theoretical claims.

change. In response to pessimistic inductionists, realist commentators have urged that the empirical success of superseded theories did not depend on their claims which conflict with present science. Rather, the theoretical constituents of these theories that were really responsible for their accurate predictions are retained in their successors, and therefore have a shot at describing the world as it is. The lesson that has been drawn from the debate over theory change then is that the realist’s epistemic optimism ought to be directed at the parts of a theory which underwrite, and explain, its predictive successes, rather than its representational content as a whole.

This idea is sufficiently vague that it can be fleshed out and interpreted in very different ways, leading to a proliferation of selective formulations of realism in the recent literature. There are many points of divergence between these variants but I will focus here on two key questions which the renormalization group sheds interesting light on: whether the selectivist should spell out their commitments in global or local terms,² and whether it is possible to implement the selective strategy prospectively, in advance of future scientific developments.

Once a broadly selective approach to the formulation problem has been adopted, the obvious question is which parts of our theories should we be committed to. There are different particular answers to this question, but there are also different types of answer. A more global answer (in my terminology) aims to provide a general characterization of the belief-worthy content of any successful theory. Saatsi (2015) calls this sort of approach ‘recipe realism’: the ideal being a formula which takes in theories and spits out beliefs about the world in a completely regular way. One brand of realism which has sometimes been understood in these terms is epistemic structural realism.³ Following Worrall (1989), contemporary incarnations of this doctrine have been inspired by episodes of theory change in physics in which posited entities, such as the luminiferous ether and gravitational field, are dropped, but continuities exist at the level of mathematical structure. The conclusion drawn is that it is the structural claims of our theories which the realist ought to put their faith in. One attempt at making this precise has been to

²The global/local axis I am invoking here should not be confused with the question of whether arguments for realism ought to be construed globally or locally. The no-miracles argument has traditionally been understood as a meta-scientific inference starting from the success of science as a whole, but this approach has recently come under attack, with some advocating a shift towards more local arguments found within science itself as the core motivation for realism (Magnus and Callender, 2004; Saatsi 2009). What I am interested in here, however, is how the realist thesis should be stated—in principle at least, this is a separate issue from how it is defended. One might favour global arguments for realism while cashing out the specific commitments engendered by your epistemology of science on a theory-by-theory basis, for instance.

³I should note here that while structural realism is often read as a global selective formulation of realism it is not clear that its advocates understand it this way. French and Ladyman (2010, 32), for instance, state that “[t]he job of predicting what will be preserved and what abandoned by future science belongs to science itself not to philosophy”, apparently disavowing the idea that the notion of structure delineates which parts of successful theories we should trust in Stanford’s (2006) sense. Still, this caricature of structural realism will be a useful foil for the localised approach to the formulation problem I will introduce shortly.

identify the structural content of a theory with its Ramsey sentence, apparently furnishing a procedure for picking out the claims the structuralist ought to commit themselves to which can be applied across the board. Many contemporary brands of realism at least aspire to this sort of blanket statement about which aspects of our theories get things right.

No formulation of this kind has achieved anything close to widespread acceptance however. One problem is that the diversity of science makes generalization a risky business. Recipes for identifying veridical content that are compelling in one area of science may be much less so in others—Stanford (2003) attacks structural realism, for instance, by pointing to examples from biology in which mathematical structure does not seem to be preserved through theory change. Another worry is that the drive towards generality leads global formulations to abstract away from peculiarities of particular theories which are relevant to appraising their representational success. While both Newtonian gravity theory and thermodynamics can arguably be described as sharing structure with more fundamental physical theories, the specific structural claims that are retained at the fundamental level are quite different in each case, as is the sort of explanation this affords of the theory's empirical successes. Even in contexts where the structuralist intuition has some purchase then, the distinction between structural and non-structural claims seems to be too coarse-grained to pinpoint the parts of a theory which drive its success. Ultimately, global formulations tend to find themselves in the uncomfortable position of being simultaneously too general and not general enough to convincingly carry out the selective project.

These sorts of concerns have led some philosophers to move towards a more local response to the formulation problem (Barrett, 2008; Saatsi 2015, 2016). Instead of trying to specify how successful theories relate to the world in one fell swoop, this approach implements the selective strategy on theory-by-theory basis. The question of what sort of doxastic attitude we ought to take towards the theoretical claims of Newtonian gravity, for instance, is answered via a close study of the theory itself, rather than invoking some general criterion for realist commitment. Barrett (2008) points, in particular, to foundational research on geometrized formulations of Newtonian gravity (surveyed in Malament, 2012), which allows us to precisely describe how some of its extra-empirical content is embedded within general relativity, as doing the real work in spelling out the sense in which the theory is approximately true. This is a thoroughly theory-specific story and the thought is that carrying out the selective strategy in practice will typically turn on local scientific arguments and resources. The local realist still adopts a general epistemic stance towards science: they take the empirical success of our theories to be explained, at least for the most part, by the fact that they are getting something right about the world. But the task of spelling out the specific claims about the world engendered by this stance is delegated to philosophers of the specific sciences.

While this move avoids some of the difficulties plaguing global formulations however, it seems to have a serious cost when it comes to the issue of prospective

applicability. The local selectivist might be able to point to intertheoretic relations with general relativity to clarify the representational status of Newtonian gravity, but what about general relativity itself? A global selective realist will have an answer here; they will wheel out their formula for identifying belief-worthy content. But what can the localist, who eschews this kind of response, say? The cases which Barrett and Saatsi point to as exemplars for the way that the approximate truth of theories can be cashed out locally invariably turn on their embedding within more fundamental theories. Consequently, realism about our current best theories might seem to end up amounting to little more than a promissory note: the local selective realist claims that general relativity latches onto reality in a way which explains its success, but exactly how we cannot yet say.⁴ Just as Newton had no idea which parts of his theory of gravity would be preserved in contemporary gravitational physics, we seem to be in a similar epistemic situation with respect to current fundamental physics.

This sort of worry is the basis of Stanford’s (2006) “trust argument”. According to Stanford any form of realism worthy of the name must tell us which features of our theories we can trust now, not merely in retrospect. I suspect that the perceived need to respond to this challenge is a key reason why defenders of realism have sought to defend a global formulation of their doctrine that can be projected into the future. An alternative reaction to this line of attack however is to simply refuse the bait and deny that the realist needs to state their epistemic commitments prospectively. After all, the claim about general relativity sketched above clearly goes beyond the constructive empiricist’s stance towards the theory—they would, of course, remain completely agnostic about general relativity’s claims about the unobservable and deny the need for an explanation of its empirical success in terms of its extra-empirical representational success. Saatsi (2016) simply bites the bullet here and calls his local version of realism ‘minimal realism’ in recognition of the fact that some intuitions about what a realist attitude towards current scientific theories amounts to are not necessarily borne out on this formulation.

It is at this juncture in the dialectic that the renormalization group becomes interesting. The application of renormalization group methods in high energy physics provides another example of how local theoretical resources can play an important role in substantiating a selective realist reading of a theory. Crucially though, this story is prospective in character, operating in advance of developments beyond the standard model of particle physics. What this suggests is that abandoning a global formulation of realism does not necessarily mean ceding the possibility of prospective commitments entirely, and a localised realism need not be as minimal as Saatsi suggests.

⁴This obviously threatens our ability to use general relativity to identify belief-worthy parts of Newtonian gravity, so this problem is apt to flow upwards, affecting non-fundamental theories as well. Barrett and Saatsi are, of course, keenly aware of this point and offer different responses. Barrett (2008) emphasizes that the transition from Newtonian gravity to general relativity can still be understood as eliminating descriptive error, while Saatsi (2016) styles it as an exemplar that can give us a handle on how a theory’s representational success could possibly explain its empirical success.

3 The Renormalization Group and Selective Realism

The renormalization group is a widely-applicable framework for investigating the behaviour of systems at different length and energy scales. The basic strategy is to study the action of a coarse-graining transformation—an operation which takes us from an initial model to a new one that lacks some of the degrees of freedom associated with variations at small length scales/high energies but shares its large scale/low energy properties. How this procedure is implemented depends a great deal on the sort of systems one is interested in, and consequently renormalization group methods take on diverse forms in different areas of physics (and beyond). I will focus here on the application of renormalization group methods to QFT, and specifically on the momentum space approach pioneered by Kenneth Wilson (Wilson and Kogut, 1974).

This story starts with the path integral expression for the partition function, Z . This crucial quantity encapsulates basically everything there is to know about a QFT model. In particular, all of a theory’s correlation functions (vacuum expectation values of field operators at disparate space-time points) can all be derived from it. For a single field ϕ , the partition function is associated with a functional integral:

$$Z = \int \mathcal{D}\phi e^{iS[\phi]}, \quad (1)$$

where $S[\phi]$ is the action of the model, and the measure $\mathcal{D}\phi$ indicates that a sum is being taken over all possible configurations of the field. As is well known, there are grave difficulties with precisely defining this integral for a field that varies over a continuous space-time. One way around this problem is to introduce an ultraviolet cutoff Λ —an upper limit on the possible momenta of field states. A straightforward way of doing this is to place the theory on a lattice so there is a minimal distance over which the field can vary. Once this is done it is possible to give a precise meaning to the path integral. The Wilsonian renormalization group is then based on setting up a coarse-graining transformation on cutoff QFT models of this kind.

Wilson’s insight was that, instead of evaluating the whole path integral at once, we can start with the contribution due to high momentum field configurations, whose fourier transforms have support above some value μ . This part of the path integral can be computed separately and absorbed into a shift in the action. In symbols:

$$\int_{|p|\leq\mu} \mathcal{D}\phi \int_{\mu\leq|p|\leq\Lambda} \mathcal{D}\phi e^{iS} = \int_{|p|\leq\mu} \mathcal{D}\phi e^{i(S+\delta S)}. \quad (2)$$

This defines a transformation that takes us from an initial cutoff QFT model to a new one, which has a lower cutoff, and a modified dynamics, but behaves like the original (specifically, sharing its long range correlation functions). This is often informally described as ‘integrating out’ the field modes associated with variations on small length scales.




Figure 1: The renormalization group flow of scalar field theories to a surface spanned by renormalizable parameters.

We can view this transformation as inducing a ‘flow’ on a space of models, with dimensions corresponding to all possible interactions between fields. Studying this flow has proved to be a powerful source of information about the scaling properties of QFT models. The most important discovery for our purposes is the phenomenon of universality in the low energy regime. It turns out that QFT models with wildly different dynamics display very similar low energy behaviour. Consider, for the sake of concreteness, the class of scalar field theories with actions of the form:

$$S = \int d^4x \left[\frac{1}{2} (\partial_\mu \phi)^2 - \frac{1}{2} m^2 \phi^2 - \sum_{n=2}^{\infty} \frac{\lambda_{2n}}{2n!} \phi^{2n} \right]. \quad (3)$$

where m is a mass parameter and $\{\lambda_4, \lambda_6, \dots\}$ are couplings for possible interaction terms. Under repeated applications of the coarse-graining transformation the renormalization group flow of this class of theories can be shown to be attracted towards a two dimensional surface spanned by m and λ_4 (as shown in figure 1).⁵ This behaviour is believed to hold generally: while infinitely many interaction terms between a set of fields are possible the renormalization group transformation induces a flow towards a finite dimensional surface spanned by so-called renormalizable parameters—those with non-negative mass dimension. This means, in essence, that large classes of QFT models look the same at suitably large length scales.

What does all this have to do with scientific realism? The thought is that these

⁵See Polchinski (1984). Duncan (2012) provides a broader discussion of these results and their significance.

renormalization group results give us the means to develop a selective realist reading of current QFTs.

On the one hand, the renormalization group helps us identify features of QFT models which we should not take representationally seriously.⁶ Quantum electrodynamics (QED) and the other component theories of the standard model of particle physics have famously produced some of the most accurate predictions in the history of science. Much of this success takes the form of estimates of cross sections for scattering events produced in particle colliders, with the current threshold on experimentally attainable energies being at the order of 10^{13} eV. The renormalization group results just discussed reveal that many features of current QFT models do not really make a difference to these empirical successes however, in the sense that they can be varied without affecting scattering cross sections at these energy scales. For one thing, it establishes that these quantities are highly insensitive to the imposition of an ultraviolet cutoff, as well as the details of how this is done. We can also vary the dynamics of a model at the cutoff scale without affecting its predictions; adding non-renormalizable interactions to the QED action, for instance, does not undermine its empirical adequacy. What this tells us is that many of the claims QFT models make about the world at the fundamental level do not contribute to, and are not supported by, the empirical success of the standard model. Of course, we also have external reasons to doubt that QFTs describe reality at all scales: the QFT framework itself is expected to break down as the Planck scale is approached. But the renormalization group gives us a precise way of pinpointing the parts of current theories that we should disavow, or at least remain agnostic about.

On the other hand, it seems to provide us with the means to articulate positive commitments supported by the success of the standard model. The classes of QFT models which share the same low energy predictions arguably make common claims about relatively large scale, non-fundamental, aspects of the world. Giving a precise characterization of this shared content is one of the central challenges facing the sort of realist view of QFT I am proposing here, but a preliminary strategy is to point to correlation functions over distances much longer than the cutoff scale as appropriate targets for realist commitment. These quantities are preserved by the renormalization group coarse-graining transformation and encode the long distance structure of a QFT model. They are also directly connected to its successful predictions—you cannot vary the long distance correlation functions of a theory without drastically affecting its low energy scattering cross sections.⁷

⁶Williams (2017) gives a similar characterization of the role that the renormalization group can play in identifying representational artifacts, and discusses some more detailed examples of distortions induced by the imposition of a cutoff, such as Lorentz violations and mirror fermions which appear when fermionic fields are put on a lattice.

⁷A potential objection here is that the realist reading of QFT models I have sketched fails to adequately distinguish itself from empiricism. This is a *prima facie* worry because correlation functions are often characterized in operationalist terms via their connection to scattering cross sections in the physics literature. A key challenge facing this approach to QFT then is to give a robustly ontological interpretation of the correlation functions, or adopt some other characterisation of the coarse-grained theoretical claims underlying the success of current

Furthermore, in demonstrating that these large scale properties of a QFT model are insensitive to what is going on at very high energies, the renormalization group is also telling us that these features are largely independent of the details of unknown physics at currently inaccessible energy scales. We thus have reason to be confident that these features of current QFTs will be retained through future theory change, in one way or another, whatever physics beyond the standard model has in store for us.⁸

The picture that emerges from these considerations then is that QFTs enjoy a kind of coarse-grained representational success, capturing some (relatively) long distance, low energy, features of the world while distorting its fundamental structure. A potentially useful comparison here is to continuum models in fluid mechanics, which misrepresent the atomic structure of real fluids but accurately describe many of their bulk properties. This fits well with the effective field theory approach to QFT that has come to dominate high energy physics in the wake of Wilson's work on the renormalization group; at least part of what is meant when physicists characterise the standard model as an effective field theory is that it correctly describes the physics of currently probed scales, but should not be trusted at higher energies. For the aspiring scientific realist this differentiated attitude towards the content of QFT models offers a way of making precise the sense in which these theories are approximately true along selectivist lines.

4 Some Morals

We have only scratched the surface of how renormalization group methods impact on our understanding of QFT, and many aspects of the preceding discussion are controversial. Wallace (2006, 2011) and Williams (2017) advance similar (and I hope complementary) analyses of the epistemic significance of the renormalization group but Doreen Fraser (2011) takes a much more deflationary line, which conflicts with some of the claims endorsed above. There remains a great deal of work to be done in developing and defending the sort of realist view of QFT just outlined then.⁹ I want to conclude however by returning to the broader question of how scientific realism ought to be formulated.

models.

⁸The idea that renormalization theory should be understood as isolating features of current theories which are robust under future theory change is advanced in a prescient paper by Alexander Rueger (1990). One concern about the way I have developed this claim here, pressed by Ruetsche (2017), is that it seems to rest on the assumption that future theories can be situated in the space of possible theories on which the renormalization group transformation acts. We can be fairly confident that future QFTs can be treated in this way, but when it comes to quantum gravity theories the question becomes much murkier. The worry then is that these renormalization group arguments ultimately fall foul of Stanford's 'unconceived alternatives' problem for realism.

⁹See James Fraser (2017) for a more detailed discussion of this approach to the epistemology of QFT which describes avenues for future work.

What sort of general morals can be extracted from this case study? First and foremost, it offers further support for a localised response to the formulation problem. The appeal to the renormalization group framework in the discussion of the previous section exemplifies the local selective realist’s claim that local scientific resources often play a crucial role in articulating the relationship a theory bears to the world. Furthermore, we found no need for an overarching thesis about which parts of our theories get things right. The resulting analysis of the representational success of QFT models does, admittedly, chime with the intuitive picture underlying epistemic structural realism—the fundamental ontological claims of QFT models were rejected while non-fundamental, broadly structural, features are singled out for realist optimism. But the putative distinction between structural and non-structural features did no real work in identifying appropriate targets for realist commitment and ultimately adds little to the picture furnished by the renormalization group. This all suggests that we ought to abandon as misguided any attempt to provide a fully general characterization of the approximate truth of empirically successful theories.

What really makes this case significant for the broader formulation debate however, is that it does not turn on the explicit embedding of a superseded theory within a more fundamental successor. This has important implications for the issue of prospective applicability. The worry, remember, was that, without a general recipe for identifying the belief-worthy content of a theory, the local realist will only be able to make a highly tentative and provisional claim about the representational success of our current most fundamental theories. Local realists like Saatsi have basically accepted this conclusion, but reject Stanford’s assertion that giving up on explicit prospective commitments means giving up on realism entirely. The renormalization group case however, suggests that we do not need a global formulation of realism to sustain prospective commitments: local theoretical resources can also play a role in supporting judgements about which parts of present theories will be preserved through theory change. The information the renormalization group provides about the dependencies between the high and low energy properties of QFT models seems to put us in a better epistemic situation with respect to the standard model than Newton was with respect to his theory of gravity. This opens up the possibility of a localised response to Stanford’s trust argument. In some scientific contexts it may be appropriate to eschew prospective judgements entirely and adopt the sort of minimal realist position advocated by Saatsi, but where scientific arguments support it, a more full-blooded realist reading of a theory, which includes commitments about which parts of its content can be trusted to remain a part of future science, may be possible.

References

Barrett, Jeffrey, A. (2008), “Approximate Truth and Descriptive Nesting”, *Erkenntnis*, 68(2), 213-224.

- Duncan, Antony. (2012), *The Conceptual Framework of Quantum Field Theory*. Oxford University Press, Oxford.
- Fraser, James D. (2017) “Towards a Realist View of Quantum Field Theory”, forthcoming in French, Steven and Saatsi, Juha (eds) *Scientific Realism and the Quantum*, Oxford University Press.
- Fraser, Doreen. (2011), “How to Take Particle Physics Seriously: A Further Defence of Axiomatic Quantum Field Theory”, *Studies in History and Philosophy of Modern Physics* 42, 126–135.
- French, Steven. and Ladyman, James. (2010) “In Defence of Ontic Structural Realism”, in Bokulich, Alisa and Bokulich, Peter (eds), *Scientific Structuralism*, Springer.
- Magnus, P. D., and Callender, Craig. (2004), “Realist Ennui and the Base Rate Fallacy”, *Philosophy of Science*, 71(3), 320-338.
- Malament, David, B. (2012), *Topics in the Foundations of General Relativity and Newtonian Gravitation Theory*, University of Chicago Press, Chicago.
- Polchinski, Joseph. (1984), “Renormalization and Effective Lagrangians”, *Nuclear Physics B*, 231(2), 269-295.
- Rueger, Alexander. (1990), “Independence from Future Theories: A Research Strategy in Quantum Theory” In: A. Fine, M. Forbes, and L. Wessels (Eds.), *PSA 1990* (Vol. I, pp. 203–211). East Lansing: PSA.
- Ruetsche, Laura, (2017), “Renormalization Group Realism: The Assent of Pessimism” *Philosophy of Science*, this volume.
- Saatsi, Juha. (2009), “Form vs. Content-driven Arguments for Realism”, In Magnus, P. and Busch, J. (eds), *New Waves in Philosophy of Science*, Palgrave Macmillan, 8–28.
- (2015), “Replacing Recipe Realism”, *Synthese*, 1-12.
- (2016), “What is Theoretical Progress of Science?”, *Synthese*, First Online, 2016. DOI:10.1007/s11229-016-1118-9.
- Stanford, P. Kyle. (2003), “Pyrrhic Victories for Scientific Realism”, *The journal of philosophy*, 100(11), 553-572.
- (2006), *Exceeding our Grasp: Science, History, and the Problem of Unconceived Alternatives*, Oxford University Press, Oxford.
- Wallace, David. (2006), “In Defence of Naiveté: The Conceptual Status of Lagrangian Quantum Field Theory”, *Synthese*, 151, 33-80
- (2011), “Taking Particle Physics Seriously: A Critique of the Algebraic Approach to Quantum Field Theory”, *Studies in History and Philosophy of Modern Physics* 42, 116-125

Williams, Porter. (2016), "Scientific Realism Made Effective", forthcoming in British Journal for the Philosophy of Science.

Wilson, Kenneth. G., and Kogut, J. (1974), "The Renormalization Group and the ϵ Expansion". Physics Reports, 12(2), 75-199.

Worrall, John (1989), "Structural Realism: The Best of Both Worlds?", *Dialectica* 43(1-2), 99-124.