A METHOD FOR SOCIAL ONTOLOGY

Iterating Ontology and Social Research

BY

DAVE ELDER-VASS

This is an electronic version of a paper published in the *Journal of Critical Realism*. Please cite as:

Elder-Vass, D. (2007) 'A Method for Social Ontology'. *Journal of Critical Realism*, 6:2, 226-49

Abstract. How should critical realism affect the practice of social science? This paper responds to this and related questions by suggesting some methodological implications of the realist theory of emergence. Given that critical realism understands causation as the interaction of emergent causal powers, and that the theory of emergence describes the type of structural relations that underpins such powers, we can practice social ontology by seeking to identify these structural relations in the social domain. Such methods, however, cannot stand or fall purely on philosophical grounds; their validity also depends on whether they work. Hence the paper briefly illustrates the application of the method to social ontology, using examples from the theory of social institutions.

Keywords: emergence, social ontology, critical realism, social structure, social institutions

How should critical realism affect the practice of social science?[1] Does critical realism, for example, have methodological and/or theoretical consequences for the social sciences, or only ontological consequences? And is critical realism compatible with a broad methodological pluralism in the social sciences, or should it be prescriptive about method?

[1] Earlier versions of this paper were presented at the IACR conference in Tromso, August 2006, the Cambridge Realist Workshop, November 2006, and a Centre for Critical Realism seminar in London in March 2007. I should like to thank the participants in those sessions for their useful comments, including Ismael Al-Amoudi, Margaret Archer, Robert Farrell, Tony Lawson, and Steve Pratten. My thanks also go to Jamie Morgan and two anonymous referees for their stimulating comments.

This paper responds to these questions by suggesting some methodological implications of the realist theory of emergence. Given that critical realism understands causation as the interaction of emergent causal powers, and that the theory of emergence describes the types of structural relations that underpin such powers, we can practice social ontology by seeking to identify these structural relations in the social domain.[2] This paper analyses the general form of such relations, from which it derives a method designed to identify the mechanisms at work in the social domain.

The first objective of the paper, then, is to describe a usable method for the practice of social ontology, which is derived from a critical realist view of ontology.[3] I say *a* critical realist view rather than *the* critical realist view because the paper proceeds from a pre-dialectical, pre-transcendental version of critical realism, as exemplified by Roy Bhaskar's first two books,[4] and thus implicitly recognises that there are already multiple versions of critical realism in circulation.[5] Secondly, I hope to show that this method is consistent with some existing methodological recommendations by prominent critical realists, and integrates them into a useful framework, although it remains only a framework, allowing for the use of a variety of methods in particular areas. I recognise from the outset that both the ontology and its translation into this method are fallible, but the method includes a response to fallibilism in the form of iteration, and so the third objective of the paper is to provide an iteration of the realist method that may then be validated and improved in future realist work. Finally, I believe the argument presented here equips critical realism with a response to some partially justified critiques that have been advanced explicitly by sociologists such as Stephen Kemp, and related ideas that are implicit in the work of figures such as Bruno Latour.[6]

The paper will begin by briefly summarising some relevant features of emergence and its implications for the general ontology of critical realism. The next section will explain how

---

[2] The centrality of emergence and causality in critical realist ontology is explored in Dave Elder-Vass, 'Emergence and the realist account of cause', *Journal of Critical Realism* 4(2) (2005): 315–38 and Tuukka Kaidesoja, 'Exploring the concept of causal power in a critical realist tradition', *Journal for the Theory of Social Behaviour* 37(1) (2007): 63–87.

[3] The method developed here is specified rather more abstractly than most methods for *empirical* research, and like all methods it is incomplete, specifying only part of what anyone adopting it would need to do. Nevertheless, it does provide a means of operationalising an ontological approach to the social sciences that has been lacking in critical realism.

[4] Roy Bhaskar, *The Possibility of Naturalism*, 3rd edn (London: Routledge, 1998) and *A Realist Theory of Science*, 2nd edn (Hassocks: Harvester, 1978).

[5] See, for example, Kathryn Dean, Jonathan Joseph and Alan Norrie, 'Editorial: new essays in critical realism', *New Formations* 56 (2005): 7–26, and Garry Potter, 'Reopening the wound: against God and Bhaskar', *Journal of Critical Realism* 5(1) (2006): 92–109.

[6] Stephen Kemp, 'Critical realism and the limits of philosophy', *European Journal of Social Theory* 8 (2005): 171–92; Bruno Latour, *Reassembling the Social* (Oxford: Oxford University Press, 2005).

this ontological framework can be translated into a series of methodological recommendations. Such methods, however, cannot stand or fall purely on philosophical grounds; their validity also depends on whether they work. Hence the third main section will briefly illustrate the application of the method to social ontology, using examples from the theory of social institutions. Finally, the broader implications of the proposed method will be discussed. It constitutes neither a complete nor a prescriptive methodology for the social sciences. It may, for example, need adjusting to cover some unusual types of structure in the social domain, and it requires complementary techniques for retroducing theory and testing it empirically.

In Search of Ontological Rigour

I would expect critical realists to be comfortable with the claim that one of the problems of the social sciences is a lack of ontological rigour. Concepts are frequently pressed into service with loose contextual definitions, with no attempt to establish what their real referents are. In the natural sciences, it would be unthinkable to employ a concept such as 'molecule' or 'black hole' or 'chimpanzee' without attempting to understand what is the range of entities that can carry this label, what they are made of, how their parts must be structured to make such an entity, what properties and powers flow from that structure, how these entities come into existence, and how their existence is maintained. But social scientists often seem happy to employ concepts such as, for example, *discourse*, *the state*, *institutions*, *values*, *money*, *value*, and *human individuals* while ignoring some or all of these questions. There is frequently a presumption that we can usefully analyse the social role of such concepts while utterly disregarding their ontological basis. Occasionally this is even justified explicitly, for example by Peter Winch:

> To discover the motives of a puzzling action *is* to increase our understanding of that action; that is what 'understanding' means as applied to human behaviour. But this is something we in fact discover without any significant knowledge about people's physiological states; therefore our accounts of their motives can have nothing to do with their physiological states.[7]

The ontological underpinnings of whatever it is that a concept refers to, then, are often taken for granted, or explicitly denied any significance, and in the extreme case concepts may be taken to refer to nothing more than themselves, thus denying the very possibility of their referents *having* ontological underpinnings. Yet these concepts are then treated as if whatever they stand for has the capability to affect the social world.

---

[7] Peter Winch, *The Idea of a Social Science* (London: Routledge & Kegan Paul, 1958), 78.

Critical realism, by contrast, offers a generalised ontological framework that identifies such causal capabilities as the real causal powers of things or *entities*. At the most general level, this framework is taken to be equally applicable to the behaviour of the natural and the social worlds. Hence realists in the social sciences proceed on the assumption that social theories must identify causal powers or emergent properties in the social world. There is a danger, however, of mirroring the ontological weakness of non-realist approaches. This fault arises if and when we find factors we believe to be causally effective and then simply label them as *causal powers* or *emergent properties* without justifying the claims that are implicit in these labels. Such a practice would provide a realist veneer to precisely the same sort of ontological superficiality that realism ought to combat.

If we are to provide a genuinely realist basis for a more ontologically rigorous approach to the social sciences, we need to bridge the gap between the general emergentist ontology of critical realism and the practice of social theorising. The starting-point, I suggest, is the categorisation of the elements of an emergentist ontology and the relations between them so that we can then investigate how particular social phenomena map onto those elements and relations.

The Elements of Emergence[8]

The first significant element of an emergentist ontology is the *entity*. Entities are wholes composed of other entities that are their *parts* (e.g. biological organisms, composed of cells, or molecules, composed of atoms), except perhaps at some lowest level that is currently beyond the understanding of our science.

Entities may possess *real causal powers*. A causal power is the capability of an entity to have a certain sort of causal effect on the world in its own right: an effect that is something more than the effects that would be produced by the entity's parts if they were not organised into this sort of whole. Real causal powers are synonymous with *emergent properties*; as the latter term is the one used in the wider (non-critical realist) literature on emergence, I shall tend to prefer it here.

Emergent properties can be explained by *causal mechanisms*.[9] Causal mechanisms are processes that depend on interactions between the parts, interactions that only occur when

---

[8] A fuller version of this argument can be found in Elder-Vass, 'Emergence'. Note that this paper assumes what I have called a *relational* version of the theory of emergence, whereas some emergentist thinkers (notably in recent philosophy of mind) define emergence in different terms that are less compatible with critical realism; see for example the strong version of emergence described by Jaegwon Kim, 'Making sense of emergence,' *Philosophical Studies* 95(1/2) (1999): 3–36, and drawn from C. D. Broad, *The Mind and its Place in Nature* (London: Kegan Paul, 1925).

those parts are organised in the particular *relations* that constitutes them into wholes that possess this emergent property.[10] Although emergent properties, and thus real causal powers, can therefore be explained, they cannot be explained away. They exist only when the relevant type of whole exists; hence they are causal powers of this type of whole and not of its parts. This means that emergentist ontologies can resolve the problem of reductionism: they allow higher-level properties to be explained scientifically (an *explanatory reduction*), but they do not allow them to be replaced with properties of the parts in causal explanations (an *eliminative reduction*).

The variety of emergence that I have been discussing so far may be termed *relational emergence*. Relational emergence is a synchronic relationship—: it describes a particular sort of relation between a whole and its parts (the relation of being composed by those parts but also possessing properties they do not) at a moment in time. Relational emergence is to be distinguished from temporal conceptions of emergence, exemplified by lay uses of the term in which it refers to nothing more than the first appearance of some phenomenon. Relational emergence entails prior or simultaneous temporal emergence, but temporal emergence does not entail relational emergence, since there can be 'heaps' – collections of entities that do not possess emergent properties as a whole. These emerge temporally but not relationally.

Temporal emergence is nevertheless important in the explanation of relational emergence, since an entity and hence its emergent properties cannot come into existence at all without a causal history. The existence of an entity at any given point in time is always contingent; it depends on a causal history in which *morphogenetic* and *morphostatic* causal factors operate. Morphogenetic factors are those that contribute to bringing about the existence of the entity in its current form, and morphostatic factors are those that contribute to sustaining that existence over time.[11] At any time, causal factors tending to end the existence of the entity may overcome the morphostatic causes tending to sustain it, hence the contingency of the entity's existence. Similarly, morphogenetic factors may *alter* the form of an entity more subtly. In some cases such changes may be consistent with the continuation of the entity's causal powers; in others, the entity may be so changed that some of its causal powers are eliminated, or indeed enhanced, or replaced with different ones. Hence synchronic emergence is fully consistent with the possibility of change.

---

[9] This is the claim that distinguishes relational from strong theories of emergence; in the latter, it is denied that emergent properties can be explained in terms of lower-level entities or properties and the relations between them.

[10] See, for example, Walter Buckley, *Sociology and Modern Systems Theory* (Englewood Cliffs, NJ: Prentice-Hall, 1967), 42; Mario Bunge, *The Sociology–Philosophy Connection* (New Brunswick, NJ: Transaction, 1999), Ch. 2.

[11] Buckley, *Sociology*, 58.

Events, finally, are caused by (actual) interactions between the real causal powers of the entities involved. Thus they are not usually determined by a single mechanism or a single 'law' as in Carl Hempel's nomological-deductive model of causation, but they are 'multiply determined' or co-determined by a variety of interacting mechanisms, which may be attributable to entities at various levels in the hierarchy of composition.[12]

To summarise, an emergentist ontology identifies a number of *structural elements* that we would expect to find in any object of scientific enquiry: *entities*, made up of *parts* (which are themselves entities), organised by particular *relations* between the parts, and possessing *emergent properties* in virtue of these relations. In order to explain these entities, relations, and properties, we need to identify the *mechanisms* by which the parts and relations lead to the properties, the *morphogenetic causes* that bring this set of parts into this set of relations in the first place and the *morphostatic causes* that keep them so. And once we are equipped with these elements, we can go on to explain *events*, and perhaps event *regularities* or partial regularities, by showing how the emergent properties or causal powers of the entities concerned interact to co-determine actual events.

The Proposed Method

So far I have been discussing an abstract, generalised, metaphysical ontology. Realists have recognised that when we apply such general ontologies to the needs of particular disciplines or groups of disciplines, we generate domain-specific ontologies that identify the sorts of elements that populate a given domain: what Ted Benton and Ian Craib call *regional ontologies*, and Bhaskar calls *scientific ontologies*.[13] This paper is concerned with providing a method for creating such regional ontologies. In a sense the method it will discuss could be applied to creating *any* regional ontology, but it has been developed from an engagement with the specific question of constructing a *social ontology*: a regional ontology for the social sciences. As a consequence, it may contain features that have inadvertently been derived from features that are unique to the social sciences, and hence may not be generalisable to other disciplines.

Critical realists should not need to be persuaded of the value of such a method. The argument that the social world cannot be theorised or explained successfully without paying explicit attention to its ontological foundations is one of the most characteristic claims of

---

12 Bhaskar, *A Realist Theory*, 110–11.

13 Ted Benton and Ian Craib, *Philosophy of Social Science* (Basingstoke: Palgrave, 2001), 5; Roy Bhaskar, *Scientific Realism and Human Emancipation* (London: Verso, 1986), 36.

critical realism.[14] As we shall see, however, once we start to examine just *how* to develop a social ontology, we also learn that this relationship works both ways: we cannot construct a regional ontology successfully without paying explicit attention to the *theory* of the discipline(s) concerned with the region. Social ontology and social theory are inextricably interwoven.

The core of the method proposed in this paper is very simple. I have already identified the structural elements of a general emergentist ontology; to develop a regional ontology, we must *map the concepts of the discipline concerned onto this structural vocabulary*. Thus, we must identify:

- the particular types of entities that constitute the objects of the discipline;
- the parts of each type of entity, and the sets of relations between them that are required to constitute them into this type of entity;
- the emergent properties of each type of entity;
- the mechanisms through which their parts, and the characteristic relations between them, produce the emergent properties of the wholes;
- the morphogenetic causes that bring each type of entity into existence;[15]
- the morphostatic causes that sustain their existence;
- and the ways that these sorts of entities, with these properties, interact to cause the events we seek to explain in the discipline.

In a well-developed science, this might seem straightforward. In each theory within the discipline, it would be clear what the entities were, what properties they possessed, what mechanisms were responsible and how these interacted in causal histories. We could read the textbooks and pop the concepts into the relevant ontological boxes. At times parts of the natural sciences have appeared to be well developed in this sense, with sets of theories that entail plausible and locally consistent scientific ontologies. Even in the natural sciences, however, scientific revolutions – such as the discoveries of quantum theory and relativity – have shaken our understandings of the corresponding scientific ontologies and reminded us that all such ontologies are inherently fallible.[16]

Yet the social sciences are at least one step further removed from the status of 'well-developed science': they consistently lack plausible, well-defined and locally consistent

---

[14] See, for example, Margaret Archer, *Realist Social Theory: The Morphogenetic Approach* (Cambridge: Cambridge University Press, 1995); Bhaskar, *The Possibility*; Tony Lawson, *Economics and Reality* (London: Routledge, 1997).

[15] The identification of the morphogenetic and morphostatic causes that contribute to an entity's development corresponds to Archer's methodological recommendation that we develop 'analytical histories of emergence': Archer, *Realist Social Theory*, 324–8.

[16] My thanks to Jamie Morgan for pointing this out.

scientific ontologies. One of the pitfalls of the social sciences is that we may assume that they *do* have such ontologies, and accept unthinkingly the sorts of ontological categorisations that appear implicit in social theories, or even in our everyday language about the social world.

For an example of the latter, consider *money*. The word is constantly used in everyday life, and frequently in the social sciences, as if money were a thing; in our terms, a type of entity. But as soon as we start to examine the ontology of money, it becomes clear that this cannot be so. For money to be a type of entity, it would have to have a characteristic type of parts, organised in a characteristic set of relations. But coins can be money, cheques can be money, and electronic transfers generated by swiping credit cards can also be money. One response to this diversity of realisers of money might be to suggest that it represents a family of types of entity, which include coins, cheques and credit-card balances among others. A more plausible one, perhaps, is that being money is a *property* possessed by a variety of different types of entity, by virtue not only of their internal structure (although this does matter) but also of their relationship to certain social institutions.[17] This would require a refinement of the portrayal of mechanisms in the previous section, since it now appears that mechanisms may rely, not only on an entity's parts and their relations to each other, but also on the relation of the whole to other entities.[18] Alternatively, we might resolve this by concluding that money is not a property of material things such as coins but rather a property of the social institution itself (see the third and fourth iterative tests below). I do not, however, propose to solve the problem of the ontological status of money here; I use the example only to illustrate the need for real work to answer such questions.

For the social sciences, then, the task of determining what type of structural element any given concept might represent is often far from trivial. On the one hand, it can be immensely difficult. The complex interrelations between physical entities, social structures, and cultural or conceptual systems in the social world make it extremely challenging to disentangle the entities and properties involved, and there are many competing schools of thought on many of these questions. On the other, resolving these questions is fundamental to resolving the ontological confusion in which these disciplines find themselves.

Applying the Method

When it comes to applying this method in practice, there is a common theme to the techniques required, which can be summed up in a single word: iterate! This section will discuss five

[17] Social institutions themselves are defined in a variety of contradictory structural terms, as we shall see below, where they will be used as an example to illustrate the use of the method.

[18] Evaluating this option would be an example of what I call the metatheory test below, that is, of reviewing our general ontology in the light of difficulties in applying it to develop a regional ontology.

types of iteration that are potentially useful, each of them characterised by submitting our proposed ontology to a particular type of test.

In general, the need for iteration arises from the combination of two factors: uncertainty and interdependence. The ways in which we formulate our inevitably fallible views about the entities involved in the social world, for example, are inextricably intertwined with our conceptualisations of their parts, relations, properties and mechanisms. This brings us to the first type of iteration, which I shall call the *local complementarity* test. We shall generally begin a regional ontology with the initial hypothesis that some concept of interest represents a particular kind of structural element, perhaps an entity or an emergent property. We then validate this initial hypothesis partly by working out the implications of this belief for the complementary set of structural elements. If we find that it is impossible to come up with a viable characterisation of a putative entity's parts, for example, as in the case of money above, we may have to revise our ontological classification of it, and this in turn will imply a new set of understandings of its complementary structural elements. Similarly, if we postulate that some concept represents a property, we shall need to identify what entity it is that possesses the property, and in virtue of what relations between its parts. Regional ontologies are not made up of isolated entities, properties and the like; they are complementary networks of mutually consistent structural elements. Only when we have a plausible complementary *set* of such classifications, then, can we move on from this point.

A second type of iteration follows from the hierarchical nature of emergence.[19] It is all very well to postulate that a certain entity, composed of certain parts, has certain emergent properties. But those parts themselves are also entities, according to the emergentist ontology, and so we must be able to justify this implicit claim for *their* ontological status as well. This is simple enough if the parts themselves are theorised uncontroversially as entities in some neighbouring discipline. Thus we can claim that human individuals are composed of cells, and take as given the status of those cells as entities because they are clearly and uncontroversially theorised as such in the biological sciences. Ideally, all putative entities should be traceable in this way to a base composed of entities theorised by the natural sciences. Let me call this the *downward recursion* test. It is common in the social sciences, however, to ignore this requirement. Signs, for example, are decomposed in structuralist semiotics into signifiers and signifieds; but it is often assumed that signifiers and signifieds are valid entity types, without any attempt to show that they in turn can be decomposed into entities that are their parts. Any ontology of cultural or conceptual systems that rests on the

[19] There is no necessity to the sequence in which these different styles of iteration are presented. They may be performed in any order, and indeed it may be necessary to iterate the iterations, going through any given type of iteration more than once.

view that these are composed of signs or symbols is ultimately incoherent unless the parts of signs themselves can be identified as entities, with their own parts, relations, and properties explained by mechanisms.

The third type of iteration is between regional ontology and regional theory. Realist regional theory is concerned with identifying the causal mechanisms underlying emergent properties (retroduction), and with explaining how these interact to produce events of interest (retrodiction).[20] But it is impossible to identify mechanisms without identifying the property to be explained, the entity possessing it, and the characteristic set of parts and relations that underpin the mechanism. In other words, retroduction depends on filling out part of the related regional ontology. Equally, as we shall see when we come to the next type of iteration, if we want to justify the claim that properties – and thus causal powers – belong to particular entities, then we need to describe the mechanism that makes them so. The consequence is that regional ontology is also dependent on retroductive theory. Hence I call this the *retroduction* test.

At this point the proposed method requires the use of established methods for empirical research, among which there are a great many different ways of retroducing mechanisms, including both quantitative and qualitative (particularly comparative) techniques, including, for example, the identification and analysis of empirical demi-regularities.[21] Hence the method for social *ontology* advocated here is potentially consistent with many alternative methods for doing explanatory social *theory*, let alone methods that pursue different cognitive interests, such as interpretation or political critique.

When we have postulated a mechanism to explain a property, a fourth variety of iteration beckons, which I shall call the *emergence level* test. This variety of iteration is implicit in the twin problems of reduction and reification. Let me illustrate this with an example: the causal power of an organisation to dismiss an employee. Now a methodological individualist might argue that organisations are nothing more than groups of individuals[22] and hence that this causal power is really a power of the individual manager who makes or communicates the decision. Such an individualist would argue that to attribute this causal power to the organisation instead would constitute an untenable reification; whereas I would argue that the individualist's attribution of this power to the individual manager constitutes an untenable eliminative reduction.

Any attribution of causal power to a particular level of the ontological hierarchy is open to *both* of these challenges, and hence to validate our own claims regarding emergent

---

20 Lawson, *Economics*, 24, 221.

21 See ibid., 204–21.

22 For example, Anthony King, 'Against structure: a critique of morphogenetic social theory', *Sociological Review* 47(2) (1999): 199–227, although King rejects the label 'methodological individualist'.

properties and the mechanisms responsible for them, we must develop and apply criteria for identifying the levels at which properties emerge. In fact a single criterion allows us to avoid both untenable reductions and untenable reifications, and this criterion is already implicit in the general theory of emergence: a property is emergent at the level where the parts of the entity possessing it would not themselves have the property if they were not organised into this sort of whole.

Hence, for example, if we attributed the power to dismiss an employee to an individual human being occupying the role of a manager, this would be an untenable eliminative reduction. A manager could not dismiss an employee unless both were parts of an organisation of a certain kind; thus the causal power is a power of the organisation, exercised on its behalf by the manager, and not a power of the manager as an individual. We could perhaps say that this is a power of the role and not of the individual occupying it, as long as we recognise that in speaking of the role here we imply the existence of the whole organisation. For an example of the opposite case, consider that managers have the power to speak to employees. If we attributed this power to the organisation rather than the manager, this would be an untenable reification: a person need not hold any particular role in an organisation, or be part of an organisation at all, in order to have this capability, and this is not altered by attaching a different description ('manager') to the person. Hence it is a power of human individuals, not of organisations or roles in them.[23]

If any hypothesised property fails this test, then we must revise our proposed ontology to relocate the property appropriately, which will in turn have implications for our understandings of the mechanisms involved.

The fifth type of iteration arises from the fallibility of these methodological proposals themselves, and indeed from the fallibility of the ontology that underpins them. From time to time we may find that the method does not work; that the ontological structures that seem most consistent with the actual social world simply do not fit the general ontology outlined here. In such cases we must always keep open two possibilities: either we have failed to apply the method and the general ontology of emergentism correctly, or that method or ontology is flawed in some way. In other words, we must be prepared to iterate the method itself, whether by making small local adjustments, or much more radical changes. Let me call this the *metatheory* test. Thus, for example, the problem with conceptual structures discussed earlier may be resolvable, on the one hand, by identifying signifiers and signifieds as 'mental entities', ultimately composed of neurons and synapses, and hence connected up to the

[23] These two examples also illustrate another significant point: that an individual can act as *both* an individual and on behalf of the organisation, exercising causal powers belonging to both levels, in one and the same action, for example, when dismissing an employee using a speech act.

ontology of the natural world (or by finding some other way of connecting the ontology of conceptual systems back to that of natural systems), or, on the other, we may have to abandon the suggestion that all entities can ultimately be connected back to the natural world in this way, and allow for the possibility of conceptual systems that have a different kind of ontological structure.

In practice, such iteration of the metatheory – the general ontology and methodology – is unlikely to take place within the confines of an applied research project. It is more likely that difficulties in a series of projects will lead to dissatisfaction that must eventually be resolved by iterating the metatheory; but the principle remains that the relationship between metatheory and theory is a two-way one. We validate *theory* by working out its implications for the actual world in the form of hypotheses that can be tested against empirical evidence, and then revise the theory if it proves inadequate to the case. In a meta-methodological parallel, we can validate *metatheory* by working out its implications for the development of theory in the form of methodological inferences that can then be tested in the process of theorising. Just as we should be prepared to revise theory that proves inadequate, we should be prepared to revise metatheory if *it* proves inadequate. The formulation of a methodology is a key moment in this cycle of validation.

The methodological techniques outlined in this section, then, involve the pursuit of consistency between the interrelated elements of a regional ontology, and between our ontological categorisations and our understanding of the empirical world. We pursue these by iterating our analysis in up to five different but interrelated dimensions:

1. *The local complementarity test* – make a hypothesis as to the type of structural element a concept of interest represents, then validate this by working out the complementary structural elements implicit in this hypothesis. If any of the latter are empirically untenable, revise your initial hypothesis and try again.
2. *The downward recursion test* – when your hypothesis claims that a concept represents a type of entity, identify the parts of this entity, and ensure that these in turn are plausibly defined as entities, following the hierarchy of composition all the way down to a level that has a well-substantiated entity status.
3. *The retroduction test* – when your hypothesis claims that a concept represents a type of property, identify the mechanism responsible for the property, and the parts and relations on which the mechanism depends. Having done so, return to the first iteration technique;.
4. *The emergence level test* – in addition, when your hypothesis suggests a property, validate the level at which the property emerges, by applying the criterion that a property is emergent at the level where the parts of the entity possessing it would not themselves have the property if they were not organised into this sort of whole.

5. *The metatheory test* – when this method seems to lead to untenable conclusions, consider revising the method.

Social Structure

Let me now try to illustrate the proposed method by applying it to the understanding of social structure, and in particular to one variety of social structure: what sociologists refer to as social institutions. The sociological sense of institution refers to normatively- endorsed social practices, such as monogamy, or turn-taking in conversation. This is distinct from the sense of *institutions* used in some other social sciences, where it is treated as more or less synonymous with *organisations* (a different form of social structure, which I have discussed elsewhere).[24] Given the complexity of the case, this will inevitably be a somewhat constrained illustration, with the result that some elements of the method will remain unillustrated, and many important aspects of social institutions will be neglected. The argument will illustrate the sort of thinking generated by this approach, rather than taking us through the process of generating it step by step. Nevertheless, I hope that it will help to demonstrate the potential of the method.

Social structure is a concept that is used by different sociologists to refer to different types – indeed, just about every conceivable type – of structural element. Jose Lopez and John Scott, in a useful categorisation of concepts of social structure, show that there are at least three senses of *structure* in widespread use in the literature: relational structure, in which social structure is identified with social relations; institutional structure, in which structure is identified with shared norms, to which we shall return below; and embodied structure, in which social structure is identified with properties or dispositions of individual human beings.[25] Some authors see structure as corresponding to social *entities* or collectivities, such as organisations. Others explicitly distinguish between collectivities and structure, seeing structure as *properties* of collectivities: 'it is only social systems or collectivities which have

---

[24] The ontological status of social structure and social institutions in particular are discussed in more detail in Dave Elder-Vass, 'Integrating institutional, relational, and embodied structure: an emergentist perspective' (paper presented at the BSA Annual Conference, Harrogate, April 2006). The ontology of organisations is discussed in Dave Elder-Vass, 'For emergence: refining Archer's account of social structure', *Journal for the Theory of Social Behaviour* 37(1) (2007): 25–44, and 'The emergence of social structure and the question of naturalism' (paper presented at the BSA Annual Conference, York, March 2005) (www.eldervass.com). A fuller account of social structure would also consider the ontology of other types of structure, including, for example, larger social systems such as capitalism.

[25] Jose Lopez and John Scott, *Social Structure* (Buckingham: Open University Press, 2000).

structural properties'.[26] And it is quite common to identify structure with *empirical regularities* in social practices: 'rule-bound and standardized behaviour patterns'.[27] Finally, a number of more recent thinkers have developed ideas of structure as embodied,[28] with the implication that it is constituted by the *parts* of society, or indeed *properties of the parts*.

In other words, social structure is a paradigmatic example of a concept whose relationship to the structural elements of an emergentist ontology is perceived in the literature as diversely as we could imagine. Part of the reason for this diversity is that different authors are sometimes genuinely referring to quite different phenomena when they are thinking of social structure—: organisations, for example, as opposed to social institutions, or statistical distributions. But much of it arises from ontological confusion about the same phenomenon.

Social Institutions: Durkheim's Dilemma[29]

We can eliminate the first source of confusion by confining our attention to one type of phenomenon. This paper will do so by concentrating on the case of social institutions. There is much less divergence in the literature on the question of how social institutions work than there is over their ontological status. Although there may be disagreements about the details, we all seem to agree that they depend on the existence of norms – shared beliefs about appropriate behaviour in certain circumstances – which produce regularised social practices as a result of their common causal effect on the individuals sharing the belief.[30]

But what exactly *are* norms, and how could they possibly be causally effective social structures? Broadly speaking, there are two sorts of answer to this question in the literature, both of which can be traced back to Émile Durkheim: individual representations and collective representations. The first answer implies that these normative beliefs are only causally effective as items of knowledge or belief held by individual human agents. The second argues that it is not individual normative beliefs but collective ones that are causally

---

[26] Anthony Giddens, *New Rules of Sociological Method*, 2nd edn (Cambridge: Polity, 1993), 'Introduction to second edition', 7.

[27] David Jary and Julia Jary, *Collins Dictionary of Sociology* (Glasgow: HarperCollins, 2000), 302.

[28] For example, Pierre Bourdieu, *The Logic of Practice* (Cambridge: Polity, 1990).

[29] Keith Sawyer has provided a persuasive emergentist reading of Durkheim in R. K. Sawyer, 'Durkheim's dilemma: toward a sociology of emergence', *Sociological Theory* 20(2) (2002): 67–85, updated in R. K. Sawyer, *Social Emergence* (Cambridge: Cambridge University Press, 2005), Ch. 6. Sawyer sees Durkheim's emergentism as resolving the dilemma between reductionist individualism and dualism. I am not convinced that Durkheim succeeded in resolving this dilemma, although I agree he was (rightly) pursuing an emergentist solution.

[30] A thorough account of the process is provided in Todd Jones, '"We always have a beer after the meeting": how norms, customs, conventions, and the like explain behavior', *Philosophy of the Social Sciences* 36 (2006): 251–75.

effective here: that individual-level normative beliefs somehow combine to give a *collective representation*, to use Durkheim's phrase, and that it is this collective representation that is causally effective.

Both arguments, however, have apparent weaknesses. If it is individual representations that cause the enactment of social practices, then we seem to be missing an explanation of the *commonality* of practices enacted by different individuals: the very thing that makes institutions what they are. The possession of beliefs by individuals seems to pass the local complementarity test; we can, for example, see these beliefs as properties of human individuals (who are the entities possessing the property), which emerge from the relations (neural networks) between their parts (cells).[31] But the explanation of social practices as an effect of individual representations does not provide an adequate explanation of the mechanism at work and thus fails the retroduction test.

The second alternative assumes that collectives as such can *have* representations, but this is highly problematic from an ontological perspective. Individuals can have beliefs because they have minds (and because they have brains), and whichever version of the philosophy of mind we prefer, it is clear that there is some sort of mechanism connecting these beliefs to neural configurations. But groups of individuals do not have minds or brains, at least not in their own right, separately from the brains of the individual members of the group. It is therefore hard to see how they can have beliefs, at least beliefs of the group itself, separately from the beliefs held by individual members of the group.[32] Collective representations, then, seem to fail the local complementarity test.

This is what we might call *Durkheim's dilemma*: should we accept the ontological implausibility of collective representations, or the mechanismic inadequacy of individual representations? Neither, of course, is a viable alternative, but the dilemma is not easily resolved, and explanations of social institutions have regularly been impaled on one horn or the other.

---

[31] This assumes a relationally emergentist account of the mind–body problem, which remains a hotly debated issue in the philosophy of mind. The most widely supported view of the mind–body problem at the moment is *non-reductive physicalism*; on this account mental properties such as beliefs are still properties of human individuals, but it is argued that mental properties cannot be reductively explained. Non-reductive physicalism would entail a different response to the retroduction test than that offered here, and indeed would represent a challenge to the metatheoretical framework adopted here, but it would still offer a response to the local complementarity test that saw beliefs as properties of human individuals.

[32] Durkheim's emergentism has always been criticised on the grounds that it seems to attribute minds to groups: for example, George E. G. Catlin, 'Introduction to the translation', in *The Rules of Sociological Method*, Émile Durkheim (New York: Free Press, 1964), xiv; Lopez and Scott, *Social Structure*, 108–9. Anne Rawls sees this as one of a series of widely held misinterpretations of Durkheim's work in her *Epistemology and Practice: Durkheim's* The Elementary Forms of Religious Life (Cambridge: Cambridge University Press, 2004), 321–4.

Collective Intentionality

One option is to replace collective representations with some sort of partial analogue that *is* ontologically feasible. This is the strategy adopted in accounts of *collective intentionality*, which has been widely debated recently.[33] One version of this (which he ascribes to John Searle and to Raimo Tuomela) is described by Geoffrey Hodgson: 'Collective intentionality arises when an individual attributes an intention to the group in which he or she belongs while holding that intention and believing that other group members hold it too'.[34]

Such attributions are inherently problematic. Given that there is no such thing as a group mind, it would be a fallacy to attribute an intention to a group. This, of course, does not mean that individuals *cannot* or *do not* perform such fallacious attributions. They may well do so, deriving a belief about the intentions of the group from evidence suggesting that a certain intention or belief is in some way typical or representative of the *individuals* making up the group. Furthermore, they formulate these beliefs verbally and communicate them to others. Let me call these *individual representations of collective intentions*. It is perhaps our tendency to formulate and refer to such representations in everyday life that makes the concept of a social norm as a collective representation seem superficially plausible. Yet these beliefs remain resolutely individual in the ontological sense: they are analytical abstractions *about* the beliefs of others (whether the person holding the belief takes them to be about typical beliefs of other individuals in the group or about some sort of group belief) but they are formed *by* individuals and they exist only as the beliefs *of* individuals. As such, they may vary between one individual and another, even when both are part of the same community, and even when both might appear to the external observer to be following (and even endorsing) the same 'rule'.[35]

I shall suggest below that individual representations of collective intentions are indeed one part of the mechanism underlying social institutions, but it should be clear that this concept of collective intentionality is not enough to save us from Durkheim's dilemma. On this definition, collective intentionality remains an individual representation, a belief held by an individual, and we are still lacking a resolution of the retroduction and emergence level tests—: an explanation of the commonality of such beliefs, and an explanation of why we

---

[33] For brief overviews of the relevant literature, see the first few pages of Christopher McMahon, 'Shared agency and rational cooperation', *Nous* 39(2) (2005): 284–308, and also Kenneth Shockley, 'On participation and membership in discursive practices', *Philosophy of the Social Sciences* 36 (2006): 67–85, esp. 75 n.6.

[34] Geoffrey M. Hodgson, 'What are institutions?', *Journal of Economic Issues* 40(1) (2006): 1–26, esp. 5.

[35] This is one of the implications of Wittgenstein's well-known arguments concerning the indeterminacy of rules; see Ludwig Wittgenstein, *Philosophical Investigations* (Oxford: Blackwell, 1953), §§143–242.

should regard such institutions as representing a causal power of a social structure, rather than purely of human individuals.

An Emergent Solution

What is missing from such accounts, I suggest, is a collective *entity* with causal powers that emerge from the relations between its parts. The linguistic usage of *social institutions* might lead us to think of these as the entities concerned, but, in an iteration driven by the local complementarity failure of the collective representations solution, I want to suggest that they are actually *properties*: properties of social groups that I shall call *normative communities* or *norm groups*. Let me clarify the argument by discussing the case of a single social institution, which tends to produce a single social practice. For the purposes of this exposition I shall assume a rather simple sort of social institution: one in which every member of the norm group both endorses the norm concerned and is expected to observe it.

I use the term *norm* here to refer to the standard of behaviour endorsed by the institution~~;~~~~,~~ but the argument so far entails that this is itself a problematic term. Given that actual endorsement of the standard is done by individuals, they may all endorse subtly different variations of a standard, and even when endorsing equivalent versions may do so in different terms. In arguing that there is one norm being endorsed and observed, we are abstracting from this variety; indeed, we are ourselves forming an *individual representation of a collective intention*, from the external standpoint of the observing social scientist. From a purist point of view, this is just as fallacious for the observer as it is for the participant in the institution. I shall suggest, however, that it is a valid move as long as we recognise (a) that when we describe a norm we are describing a typical or representative understanding of that norm in the group, and not a collective representation as such; and (b) that there is no well-defined method for identifying 'typical' or 'representative' beliefs in a group, and hence that such identifications may be contentious, and certainly will be fallible. Despite these qualifications, we may find in practice that certain norms are so consistently understood within a group that we can neglect these problems.

One consequence of the simplifying assumption that every member of the norm group both endorses the norm concerned and is expected to observe it~~,~~ is to set aside issues of differential power within norm groups, which arise when the group endorsing the norm is different from the group expected to observe it. This is one of a number of important issues that would need to be elaborated on further in any account of institutions that was to be considered adequate in its own right, rather than simply as an illustration of an ontological method.

In this emergentist account, the institution is the property or power (of the norm group) that tends to produce the practice concerned.[36] The parts of the group are the individuals who are its members (hence, by contrast with accounts assuming collective representations, we now have an account that satisfies the downward recursion test). These individuals share broadly similar beliefs that this practice is considered appropriate in certain circumstances, and in a pragmatic sense they can validate that these beliefs are shared through their relationships with adjacent members of the community, both by asking them directly about the normative beliefs concerned, and through their experience of how these adjacent members react when the relevant behaviours occur. Hence they will tend to form similar individual representations of collective intention with regard to the norm. These similarities are essential to the emergence of the institution, since any commonality of social practice would tend to break down in the face of highly divergent expressions of belief concerning their desirability. But as I have argued above, they are not enough to take us beyond the *individual representations* horn of Durkheim's dilemma.

What we need, in order to move beyond this dilemma, is to identify a *relation* between the members of the group that makes them act differently than they would if they were not members of the group. The relation that makes a difference, I suggest, is not these shared beliefs but the commitment that the members of the group have to *endorse* and *enforce* the norm with each other, whether by advocating it, by praising or rewarding those who enact it, by criticising or punishing those who fail to enact it, or even just by ostentatiously enacting it themselves.[37] The consequence of such endorsement is that the members of the group know (whether consciously or subconsciously) that they face a systematic incentive to conform to the norm. Not only will other individual members of the group take an incentivising stance, but when they do so they will be taken to be acting on behalf of the group as a whole and will be supported by other members of the group.[38] Thus we now have a clear picture of the

---

[36] Like all causal powers in the critical realist model, norms do not *determine* behaviour but *contribute* causally to its determination, alongside other causal factors, and hence they only *tend* to produce a given outcome. See Bhaskar, *A Realist Theory*, and Elder-Vass, 'Emergence'.

[37] The role given here to endorsing and enforcing *practices* in the explanation of social institution is reminiscent of Rawls's solution in *Epistemology* to Durkheim's dilemma. But while ethnomethodologists such as Rawls practice a variety of central conflation in substituting practices for both individuals and social structures, this paper sees them as part of a mechanism that sustains social structures while depending on individuals.

[38] These incentives may be normalised for the members of the group, in the sense that they come to regard them as moral standards they observe independently of a structure of incentives; we could in principle test whether the incentive structure provided by mutual interaction remains important by removing it. If the incentive structure *is* significant, adherence to the norm would tend to fade as people realised that the incentivising behaviours of the other members of the norm group had changed.

mechanism, in terms of how it arises from the relations between the parts of the entity possessing the property, and the retroduction test is satisfied.

As a consequence of being members of a norm group, then, these individuals act differently than they would if they were not such members. Even if they held the same normative belief, they would not necessarily act in the same ways regarding it (either endorsing it so strongly or enacting it so frequently) if they were not part of a group that shares a commitment to endorse and observe the norm.[39] And it is not the mere aggregation of other people's beliefs that makes the norm group effective; it is the commitment to mutual *interaction* between the members in support of the norm, and the knowledge that they can rely on the support of others in such action, that makes the norm group more effective in enforcing the norm than the same number of individuals would be if they did not share this mutual commitment. These *relations*, then, when combined with these sorts of *parts*, provide a generative *mechanism* that gives the norm group an emergent *property* or causal power—; thus, we have satisfied the local complementarity test. This property or power is the institution, which tends to affect the behaviour of members of the norm group, and by demonstrating why it is at *this* particular level that the property arises, we have also satisfied the emergence level test.

Social Institutions: Socially Real or Materially Real?

One consequence of this argument is to clarify the relation of social structure to material reality. Steve Fleetwood has argued, in an otherwise admirable paper, that social structures are 'socially real' rather than 'materially real' entities. 'Socially real entities', he argues, 'have the following properties … they contain not one iota of materiality, physicality, solidity or whatever. We cannot touch a social entity'.[40]

But my argument here leads to a rather different conclusion. Social institutions, on this account, are properties of groups of people. As entities composed of material entities – people – norm groups are themselves material. Granted, there are some interesting differences in the sorts of relations that obtain between the people making up a norm group and those

---

[39] As Shockley, points out, one emerging point of agreement in the literature on collective intentionality is that 'membership in a group can serve to provide or generate reasons for individual actions, reasons which would not be in place without membership in that group' ('On participation', 75 n.6).

[40] Steve Fleetwood, 'Ontology in organization and management studies: a critical realist perspective,' *Organization* 12 (2005): 197–222, esp. 201. His non-material 'ideationally real' entities are also problematic, although that is beyond the scope of this paper.

between the parts of most natural entities. Let us briefly consider two of the most significant.[41]

First, the parts of a norm group may be distributed rather widely and unstably in space, with the parts moving around independently of each other and with many other material entities in the intervening spaces, whereas we tend to think of natural entities as being composed of parts in tightly constrained adjacent spatial relations to each other. While this may be true of many natural entities – animals and plants, atoms and molecules, for example – even natural entities may be rather more fluid in form. Consider rivers, for example; their parts are constantly changing, constantly moving relative to each other, with entities such as rocks in the riverbed, boat hulls and fishes intervening in their space, and even their overall shape may change, for example when rains swell the size of the river or meanders are cut off to form ox-bow lakes. Norm groups take this fluidity of form one step further, dispensing with the adjacency of parts that still remains to fix the spatial relations of entities like rivers, but why would this lead us to believe that such groups are any less *material* than more adjacent forms, such as a human pyramid?

Secondly, social collectivities have non-exclusive parts: a member of one norm group may also be a member of many others, and also of other social structures such as organisations. By virtue of their overlapping membership of multiple social entities, human beings are subject to (possibly conflicting or contradictory) influences from a variety of different organisations and institutions in which they are expected to play a role.[42] By contrast, we usually think of the parts of a natural entity as being exclusive to that entity, or at least to one and only one entity at a given level of the ontological hierarchy. A molecule, for example, can be part of only one cell at any given moment, although even in nature~~,~~ an entity can be part of a variety of nested entities at different levels: the molecule may also be part of an organism, of which the cell concerned is a part. To visualise the materiality of norm groups, we must envisage not only the nesting of entities but also the intersection of multiple entities at the same level. These differences might distract us, but they do not alter the essential point: that norm groups are composed of parts that are material entities, and therefore are material entities themselves.

Social institutions, then, are properties of material entities. Now one might argue that properties cannot be material, nor, indeed, non-material. We might say that only entities can be material or non-material: that properties are inherently *amaterial*. But this, of course,

---

41 I have discussed the implications of these differences at more length in Elder-Vass, 'The emergence of social structure and the question of naturalism' (www.eldervass.com).

42 Hence their performance of one role may sometimes be affected by the influence of the others; see Erving Goffman, *The Presentation of Self in Everyday Life* (Edinburgh: University of Edinburgh, Social Sciences Research Centre, 1956).

would be equally true of the properties of other material entities. Being green, or having a certain mass, on this account, would also be amaterial properties. Alternatively, we might define the materiality of properties by reference to the materiality of the entities possessing them. On this account, social institutions, like being green or having a certain mass, would be material properties. On either account, however, social institutions are no less material than the properties of natural physical entities; and the *entities* in this regional ontology – the norm groups themselves – are just as material as any natural entity.

Conclusion

This paper, then, has outlined and illustrated a method for developing a regional ontology for the social sciences, a method that is explicitly derived from the general or metaphysical ontology of critical realism, understood in emergentist terms. This general ontology entails that the world is populated with certain types of structural elements, in certain types of relation with each other, and the proposed method identifies these types of structural element among the phenomena investigated by the social sciences.

Where sciences are well developed such an exercise will not tell us anything new; but the social sciences are riddled with failures of ontological clarity, and in these circumstances the method is potentially very useful. Arbitrary assignations of phenomena to structural categories, however, do not offer a significant improvement on the status quo; we must be able to validate such proposals if they are to be robust and useful. This paper therefore proposes a series of iterative tests, again derived explicitly from the general ontology, which can be used to validate, and if necessary reject, such proposals. An interlinked set of structural elements identified and validated using this procedure is likely to constitute a valuable resource for theory development in the discipline concerned, and indeed can only be developed in concert with a process of theory development.

This argument has been illustrated with the case of social institutions, one of the types of social structure. I hope that this illustration has begun to demonstrate the value of the proposed method, showing that it can take us beyond Durkheim's dilemma, beyond alternatives such as collective intentionality, and beyond claims of the non-materiality of the social, to a potentially more productive understanding of social institutions. Nevertheless, like all methods its value can only be confirmed by repeated application, and the learning from experience that results. Hence I have emphasised the potential need to iterate the method itself (and even the ontology on which it is based) as well as the particular hypotheses that are developed when we apply it.

While some critical realists may be concerned by the implication that both realist accounts of social structure and its canonical accounts of ontology should be open to revision,

such openness ultimately holds out the possibility of a stronger, not a weaker, ontology. Furthermore, it constitutes a strong response to a number of critics. Thus, for example, Latour has criticised what he calls 'sociology of the social' for making assumptions much too readily about what can be counted as a social structure;[43] by contrast, this method asks us to question those assumptions wherever we find them, and provides a way for us to validate or reject them.

Secondly, Kemp has criticised the one-sided relationship that critical realists seem to advocate between philosophy and social science, arguing that 'social scientific research should be conducted without philosophical legislation',[44] and that 'it will not be philosophy disconnected from empirical research that leads the social sciences towards success'.[45] This paper replaces the kind of one-sided relationship criticised by Kemp with a two-way relationship that does use philosophy to inform the research process, but also allows philosophy to learn from that process in return. This is, after all, very much in the spirit of Bhaskar's *A Realist Theory of Science*, which arguably begins this particular style of iteration by deriving critical realism's ontology from the process of scientific research.

Nevertheless, the implication of this paper is that the social sciences *can* benefit from the application of critical realism's ontology. What a critical realist social science can offer that is both characteristically different from existing approaches, and decisively linked to the realist philosophy, is the recognition of the importance of regional ontologies for theory-building. What this paper has sought to provide is a characteristically realist method for developing such ontologies. This does not purport to supplant existing methods for social research; it is *a* method for social ontology, rather than *the* method for social science. And it is itself a moment in a process of iteration; it may come to be refined or indeed rejected on the basis of our experience of working with it. But it does, I claim, offer a potentially important addition to the set of methods available to realist social scientists.

Bibliography

Archer, M. *Realist Social Theory: The Morphogenetic Approach*. Cambridge: Cambridge University Press, 1995.

Benton, T. and I. Craib. *Philosophy of Social Science*. Basingstoke: Palgrave, 2001.

Bhaskar, R. *The Possibility of Naturalism*, 3rd edn. London: Routledge, 1998.

Bhaskar, R. *A Realist Theory of Science*, 2nd edn. Hassocks: Harvester, 1978.

43 Latour, *Reassembling*.

44 Kemp, 'Critical realism', 171.

45 Ibid., 187.

Bhaskar, R. *Scientific Realism and Human Emancipation*. London: Verso, 1986.

Bourdieu, P. *The Logic of Practice*. Cambridge: Polity, 1990.

Broad, C. D. *The Mind and its Place in Nature*. London: Kegan Paul, 1925.

Buckley, W. *Sociology and Modern Systems Theory*. Englewood Cliffs, NJ: Prentice-Hall, 1967.

Bunge, M. *The Sociology–Philosophy Connection*. New Brunswick, NJ: Transaction, 1999.

Catlin, G. E. G. 'Introduction to the translation'. In *The Rules of Sociological Method*, É. Durkheim. New York: Free Press, 1964.

Dean, K., J. Joseph and A. Norrie. 'Editorial: new essays in critical realism'. *New Formations* 56 (2005): 7–26.

Elder-Vass, D. 'Emergence and the realist account of cause'. *Journal of Critical Realism* 4(2) (2005): 315–38.

Elder-Vass, D. 'The emergence of social structure and the question of naturalism'. Paper presented at the BSA Annual Conference, York, March 2005.

Elder-Vass, D. 'For emergence: refining Archer's account of social structure'. *Journal for the Theory of Social Behaviour* 37(1) (2007): 25–44.

Elder-Vass, D. 'Integrating institutional, relational, and embodied structure: an emergentist perspective'. Paper presented at the BSA Annual Conference, Harrogate, April 2006.

Fleetwood, S. 'Ontology in organization and management studies: a critical realist perspective'. *Organization* 12 (2005): 197–222.

Giddens, A. *New Rules of Sociological Method*, 2nd edn. Cambridge: Polity, 1993.

Goffman, E. *The Presentation of Self in Everyday Life*. Edinburgh: University of Edinburgh, Social Sciences Research Centre, 1956.

Hodgson, G. M. 'What are institutions?' *Journal of Economic Issues* 40(1) (2006): 1–26.

Jary, D. and J. Jary. *Collins Dictionary of Sociology*. Glasgow: HarperCollins, 2000.

Jones, T. '"We always have a beer after the meeting": how norms, customs, conventions, and the like explain behavior'. *Philosophy of the Social Sciences* 36 (2006): 251–75.

Kaidesoja, T. 'Exploring the concept of causal power in a critical realist tradition'. *Journal for the Theory of Social Behaviour* 37(1) (2007): 63–87.

Kemp, S. 'Critical realism and the limits of philosophy'. *European Journal of Social Theory* 8 (2005): 171–92.

Kim, J. 'Making sense of emergence'. *Philosophical Studies* 95(1/2) (1999): 3–36.

King, A. 'Against structure: a critique of morphogenetic social theory'. *Sociological Review* 47(2) (1999): 199–227.

Latour, B. *Reassembling the Social*. Oxford: Oxford University Press, 2005.

Lawson, T. *Economics and Reality*. London: Routledge, 1997.

Lopez, J. and J. Scott. *Social Structure*. Buckingham: Open University Press, 2000.

McMahon, C. ‘Shared agency and rational cooperation’. *Nous* 39(2) (2005): 284–308.

Potter, G. ‘Reopening the wound: against God and Bhaskar’. *Journal of Critical Realism* 5(1) (2006): 92–109.

Rawls, A. W. *Epistemology and Practice: Durkheim's* The Elementary Forms of Religious Life. Cambridge: Cambridge University Press, 2004.

Sawyer, R. K. ‘Durkheim's dilemma: toward a sociology of emergence’. *Sociological Theory* 20(2) (2002): 227–47.

Sawyer, R. K. *Social Emergence*. Cambridge: Cambridge University Press, 2005.

Shockley, K. ‘On participation and membership in discursive practices’. *Philosophy of the Social Sciences* 36 (2006): 67–85.

Winch, P. *The Idea of a Social Science*. London: Routledge & Kegan Paul, 1958.

Wittgenstein, L. *Philosophical Investigations*. Oxford: Blackwell, 1953.