RE-EXAMINING BHASKAR'S THREE ONTOLOGICAL DOMAINS: THE LESSONS FROM EMERGENCE

Dave Elder-Vass

[This is a pre-publication version of a chapter published in an edited collection. Please cite as:

Elder-Vass, D. (2006) 'Re-examining Bhaskar's three ontological domains: the lessons from emergence'. In C. Lawson, J. Latsis & N. Martins (Eds.), *Contributions to Social Ontology*, 160-76. London: Routledge.]

Introduction[1]

Although Roy Bhaskar's ontology in *A Realist Theory of Science* (Bhaskar, 1978) is explicitly stratified into a hierarchy of levels, he makes relatively little attempt to examine the basis of this stratification, or its implications for his three domains of the empirical, the actual, and the real (Collier, 1994, p. 130).[2] This paper attempts to remedy that deficiency, by investigating the nature of a stratified reality based on emergence, and considering how this impacts our understanding of experiences, events, entities, and causes. Bhaskar uses 'stratification' to indicate two quite different ontological schemes: the stratification of the world into emergent explanatory levels, which I shall call *level stratification*, and the division of ontology into domains, which I shall call *domain stratification*. It is the relation between these two different schemes that is the central theme of this paper.

I hasten to emphasise that my objective here is primarily to refine Bhaskar's argument, and to repackage it in a form that provides greater clarity, rather than to undermine or contest its essential content. My argument in no way conflicts, for example, with claims for the existence of level stratification, or with the need to separate causal powers from actual causation and both from empirical experience. What the paper does seek to do, on the other hand, is to add some depth to the characterisation of experiences, events, and entities, and to examine their relationship to Bhaskar's ontological domains. Ultimately this will lead it to question the nature of

[1] I would like to thank Jason Edwards, the participants at the IACR 2004 conference, and three anonymous referees for their useful comments on earlier drafts of this paper.

[2] I assume the reader has a certain degree of familiarity with some of Bhaskar's concepts, notably the transitive/intransitive distinction.

Bhaskar's distinction between the domains of the actual and the real, and to consider an alternative way of looking at this distinction.

The paper will begin by introducing the key terms: Bhaskar's three domains, the concept of emergence, and two different ways of looking at multi-layered entities and events. It will then move on in turn to discuss the implications of level stratification for events, entities, causes, and experiences. Finally it will bring together the threads of the argument to re-evaluate Bhaskar's three domains.

Bhaskar's three domains

In *A Realist Theory of Science*, Bhaskar argues from the intelligibility of experimental activity to the conclusion that "there is an *ontological* distinction between scientific laws and patterns of events" (Bhaskar, 1978, p. 12). Such laws depend upon the existence of 'natural mechanisms', and "it is only if we make the assumption of the real independence of such mechanisms from the events they generate that we are justified in assuming that they endure and go on acting in their normal way outside the experimentally closed conditions that enable us to empirically identify them" (p. 13). Similarly, "events must occur independently of the experiences in which they are apprehended. Structures and mechanisms then are real and distinct from the patterns of events that they generate; just as events are real and distinct form the experiences in which they are apprehended. Mechanisms, events and experiences thus constitute three overlapping domains of reality, viz. the domains of the *real*, the *actual*, and the *empirical*" (p. 56). The relationship between these domains is summarised in a table, reproduced below as Figure 1.

	Domain of Real	*Domain of Actual*	*Domain of Empirical*
Mechanisms	x		
Events	x	x	
Experiences	x	x	x

Figure 1 – Bhaskar's three domains: populating entities (Bhaskar, 1978, p. 13)

Bhaskar clearly intends the domain of the empirical to be a subset of the domain of the actual, which in turn is a subset of the domain of the real (Bhaskar, 1978, Note to Table 1, p. 56; Bhaskar, 1993, p. 207). We can represent these inclusion relations in a Venn diagram (see Figure 2).

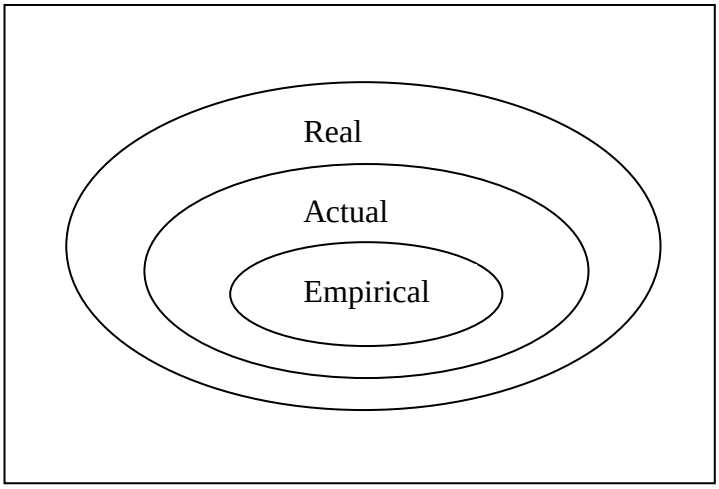

Figure 2 – Bhaskar's three domains: inclusion relations

Emergence and its basis

The second element of Bhaskar's ontology with which this paper will engage is the stratification of the intransitive world into levels – the atomic, the molecular, the biological, and the like. This level stratification depends upon the phenomenon of emergence, which is most simply described as the relationship which makes it possible for a whole to be more than the sum of its parts. Bhaskar himself defines emergence as "the relationship between two terms such that one diachronically, or perhaps synchronically, arises out of the other, but is capable of reacting back on the first and is in any event causally and taxonomically irreducible to it, as society is to nature or mind to matter" (Bhaskar, 1994, p. 73).

To put it more simply, emergence occurs when a whole has properties or powers that are not possessed by its parts. In this sense, the concept of emergence is inherently compositional, in the sense that higher-level entities always emerge from collections of lower-level entities that are their components or parts.[3] [4]

But how is emergence possible? If we accept that emergent wholes have properties that are not possessed by their parts, then where do those properties come from?

Like many others, I argue that they come from the organisation of the parts, from the maintenance of a stable set of relations between the parts that constitute them into a particular kind of whole (Archer, 1982, p. 475; Cilliers, 1998, p. 43; Cunningham, 2001, p. S68; Emmeche et al., 1997, p. 106; Lloyd Morgan, 1923, p. 64; Sawyer, 2001, p. 560; Sperry, 1986, p. 266).[5] Thus it is the fact that a higher-level entity is composed of a *particular stable organisation* of lower-level entities that

[3] Bhaskar sometimes uses 'higher' and 'lower' in the opposite sense to this, but I shall maintain the usage that 'lower' entities are components of 'higher' entities for the sake of consistency with most other work on emergence.

[4] Although sometimes critical realists seem to adopt a compositional definition of emergence (see, for example, (Collier, 1998, p. 264)) at other times this is denied (e.g. (Bhaskar, 1978, p. 169; Collier, 1994, p. 116)). I defend the compositional approach against Collier's argument in (Elder-Vass, 2005). The compositional approach is implicit in virtually all of the non-critical-realist literature cited in this paper.

[5] Useful histories of the concept of emergence can be found in (McLaughlin, 1992) and (Blitz, 1992).

gives it the possibility of exerting causal influence in its own right. In other words, it is the set of relations between the lower-level entities that makes them 'more than the sum of their parts'. Only when this particular kind of parts is present in this particular set of relations to each other does the higher level entity exist, and only when this particular kind of parts is present in this particular set of relations to each other do they have the causal impacts that are characteristic of the higher-level entity. As Archer puts it, "Emergent properties are therefore relational: they are not contained in the elements themselves, but could not exist apart from them" (Archer, 1982, p. 475). Note that a higher-level entity is only emergent when it just so happens that, when a set of lower level entities is so organised as to create it, the resulting entity has a consistent causal impact that is not a simple summation of the impacts of the its components. Now, the particular causal influences that any particular entity type may exert, and the way in which the presence of its parts in the required relations produce these higher level effects, are a matter for the particular science of the case – we cannot go any further at the philosophical level in explaining why particular cases of emergence work.

We can go further, however, in identifying another general pre-requisite for emergence. As the existence of the whole is inseparable from the continuing presence of the required parts in the required arrangement, then emergence itself depends upon the set of causes that maintain a set of such parts in just such an arrangement. Bhaskar has commented on the dual aspect of emergence as a synchronic and diachronic relation; but the diachronic aspect of his account seems to relate to the original creation of the new level of reality. While this original creation is clearly necessary, the maintenance of the particular entities that constitute that new level is equally important. There is not only a set of causes that brings the entity about, but also a further (possibly overlapping) set that maintains its continuing existence – what I will call, after Buckley, its *morphostatic causes* (Archer, 1982, p. 480, n8; Buckley, 1998, p. 53). It is these causes that are responsible for the stability or persistence of the entity.

The role of these morphostatic explanations of continuity of structure is critical to emergence. Any number of accidental combinations of lower-level entities may be brought about by a vast range of morphogenetic causes over the course of time, but it is only those combinations that have continuity of structure that persist. At any time, it is possible that a more powerful morphogenetic cause may overcome the morphostatic causes for any given entity. At this point, the emergence of the higher level entity is dissolved. It is the ability of morphostatic causes to resist such effects that sustains the existence of higher-level entities and hence any emergent properties they may have. But this continuing existence is always dependent upon the uncertain outcome of the ongoing tension between these different types of cause.[6]

Emergence, then, is the outcome of a process by which a set of morphostatic causes, which may be both internal and external, sustain a set of lower level entities in relationships that constitute them into a stably organised higher level entity that can as a result exercise powers that are not possessed by its component entities either in isolation or in an unstructured aggregation.

[6] There is a useful role in the explanation of morphostasis for concepts like 'negative feedback' from cybernetics and 'strange attractors' from complexity theory. Some realists have sought to explain morphostasis in terms of the necessity of 'internal relations', e.g. (Archer, 1995, p. 173; Sayer, 1992, p. 119). This is useful where it is taken to mean that certain combinations of parts tend by *natural* necessity to hold together; but not when it is read as a claim for the *logical* necessity of certain relations. I intend to discuss this question in more depth elsewhere.

Level abstracted and downwardly inclusive views

One implication of emergence is that entities with emergent properties or powers are themselves composed of other such entities, which are in turn so composed, and so on.[7] A plant, for example, consists of cells, the cells consist of molecules, the molecules consist of atoms, and so on. Any given entity, then, can be seen as internally stratified into many different levels or layers, each level representing sets of parts that are combined into the entities at the next level up.

Now, for most purposes, when we discuss any given entity we are in the habit of ignoring the role of its parts. To treat an entity in this way is to take what I propose to call a *level abstracted* view of it – i.e. a view that considers the effects of the whole entity in isolation from the existence or effects of its parts. I will argue in this paper, however, that for other purposes we sometimes need to treat a whole entity quite explicitly as a stratified ensemble of parts at various ontological levels. This is to take what I call a *downwardly inclusive* view of the entity. These two terms are illustrated in Figure 3 below.

Here, L1 represents the highest level of a whole – to continue the example, a plant. L2 represents the first decomposition of the whole into its parts – in this case, perhaps, the cells of the plant, and the relations between them that constitute them into a whole plant. L3 represents the next decomposition – here, the molecules that make up the cells and the relevant relations between them. And the pyramid may continue downwards, until its base is lost in the mists of our limited understanding of sub-quantum science. Of course, a plant is not made up of the whole plant *plus* its cells *plus* its molecules, and so on; each of these levels represents a different decomposition of the same whole; it is only our view of the plant that must sometimes encompass the recognition that the whole plant is simultaneously each of these different decompositions.

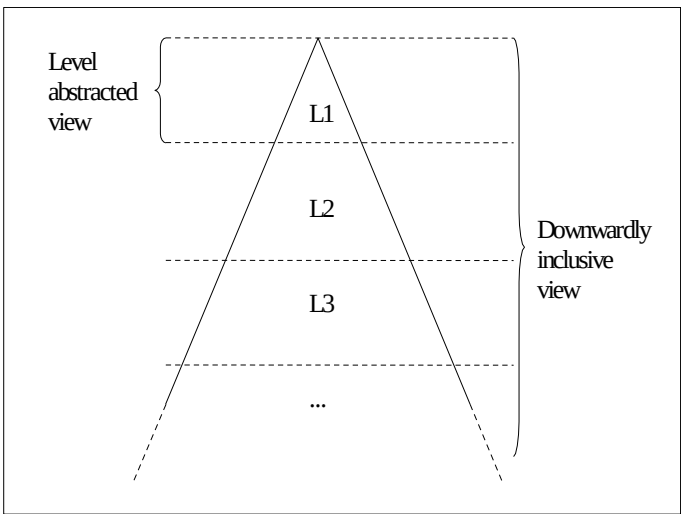

Fig 3 – Internal stratification

[7] It is not clear in the current state of science whether this nesting proceeds indefinitely or whether there is some lowest level of entity that will eventually be reached in this series of progressive decompositions. We can ignore this question for the purposes of the argument presented here.

With these preliminaries completed, we can now turn to examining the nature of the elements that inhabit Bhaskar's three domains in a level stratified world.

Experiences

Let me begin with experiences. There are two key factors that influence the shape of our experiences.

First, as Bhaskar tells us, "Experiences, and the facts they ground, are social products" (Bhaskar, 1978, p. 57). Experiences are social products because our experiences are not simply a set of sense-data, but rather the result of our application of a socially-influenced conceptual framework to the interpretation of that sense-data. Our eyes may detect a pattern of colours; but what we experience is 'seeing' a set of meaningful objects behaving in meaningful ways. It is in this interpolation of our conceptual frameworks between sense data and 'experience' that experiences become 'social products'. Hence 'experiences' are no longer purely the outcome of the events they might appear to reflect, but rather a product of the combination of those events with our prior knowledge.

Second, our experiences, despite being interpreted, are constructed on the basis of our sense-perceptions. Those sense perceptions are inevitably limited to impressions of those segments of reality that we are capable of perceiving with the senses we possess, as augmented by any artificial tools that are available.

Now, the combined effect of this process of interpretation and our restricted perceptual (and perhaps cognitive) abilities is that we generally perceive reality as 'flat' in the sense that our experiences are interpreted as impressions of entities at a single level of stratification. When we perceive the human being, we do not simultaneously and inseparably perceive the organs, the cells, the molecules that make them up. If we perceive the cells of a living tissue under a microscope, we do not simultaneously and inseparably perceive the organ or the organism to which they belong.

Thus, our experiences are already, through the process of abstraction that is inherent to perception, and as a result of the limited slice of reality to which our senses give us access, level abstracted views of what is in actuality an inherently multi-levelled occurrence.

Events

Now, because of the nature of our experiences, our everyday (empirical) concept of an 'event', which we take to be the naturally-given subject of any explanation in science, is itself an abstraction from reality. Thus, when we say, for example, 'the pen fell on the floor', we are already, in framing our reportage of the event, making an assumption about which abstraction from what was happening in a multi-level stream of interconnected happenings is the one that is relevant and requires explanation. We could have looked at the same happenings and chosen to explain the behaviour of the molecules or atoms involved, or the writing process or the world-historical events or the social history of which the falling of the pen formed a part. But in selecting out one of these happenings from the rest as the thing to be explained, we have already created the illusion that this is an event that can be given an explanation in its own right, independently of its component events and of the larger events of which it forms a part: here we have a *level abstracted view of the event*.

In seeing events as level abstracted, we implicitly frame the (retroductive) question of how they are caused in a way that demands explanations of a particular form - in a way that pushes us into thinking in terms only of a particular stratum or level of organisation.[8] But any causal account of an event seen in level abstracted terms forms only part of a larger picture. A more complete explanation can always be provided by re-integrating the event into the larger stratified picture of which it forms a part, and relating its explanation to the explanations of the other events in which it is inextricably implicated, either as subset or superset.

Now, I suggest that the way to make sense of causal explanations of individual events in this context, where an event is defined as the behaviour of a given entity at a given time, is to allow that every event inescapably includes the behaviour of the composing lower level entities as well (Lloyd Morgan, 1923, p. 15). To view an event in these terms is to see it in *downwardly inclusive* terms. It might seem that we could also look at the event in *upwardly inclusive terms*, in which sense it would also include the behaviour of all the higher level entities of which the first entity is a part. But this seems inherently infeasible, given the indeterminate (and indeterminately large) range of higher level entities that may be part of this set, all the way up to the universe itself. There is no apparent reason why our interest in the falling of a pen, for example, should also require us to be interested in that complete set of higher level events, even for metaphysical purposes. We may, of course, be interested in some *particular* higher level event of which the falling of the pen is part, but if that is so we can take a downwardly inclusive view of that higher level event, which will include the behaviour of the pen. As a general rule, then, we need not take an upwardly inclusive view of an entity or event.

Hence, in explaining a downwardly-inclusively-conceived event, we recognise, for example, that when a pen drops, it is inseparably part of this event that the components composing the pen remain in a set of relationships through which they constitute the pen, and behave in whatever ways are required for the pen to drop. This is the inevitable consequence of the set of morphostatic causes whose operation must be present for the pen to exist as such through the entire course of this event. Thus, the various material parts of the pen go through a set of events that forms part of the higher (downwardly-inclusively-conceived) event, the molecules that compose those parts go through another set that also forms part of the higher events, and so on through the atoms, subatomic particles, and so forth. Given that we do not have fully adequate understandings of the lower end of this spectrum, we must accept that only partial descriptions and hence only partial explanations are possible of the lower-level set of events that composes the higher level event. For most practical purposes we can and indeed must ignore the lower levels of this hierarchy, but for the purpose of understanding the ontology of events and causation we must recognise their significance.

If we wish to understand the role of emergence in individual events, and the relations between causes at different emergent levels, then, the correct (downwardly inclusive) account of individual events is inherently level stratified. We need to recognise that the events which populate Bhaskar's 'domain of the actual' are downwardly inclusive and multi-levelled. This clearly corresponds to Bhaskar's conception of the actual as that domain of reality in which a vast range of particular causes interact to cause events. And on my account the actual includes not only events that are unobserved by virtue of the absence of an observer, but also those levels of

[8] On retroduction, see (Lawson, 1997, p. 24).

multi-levelled events that are unobserved by virtue of operating below (or above) the perceived levels of reality.

Entities

Now, events involve the behaviour of things, or entities. Like events, entities are inherently and inclusively multi-levelled, but when we label them in empirical experience, and also when we employ them in causal statements, we typically abstract from most of those levels.

Where do entities fit in Bhaskar's three domains? If events are constituted by the behaviour of entities, and if (downwardly inclusive) events belong in the domain of the actual, then it would seem clear that (downwardly inclusive) entities must also belong in the actual. Indeed, at one point Bhaskar indicates that "the domain of actualities… may be extended to include things as well as events" (Bhaskar, 1978, p. 32). The claim that what we experience is a subset of the actual would also seem to support this argument – clearly we can experience things as well as events,[9] and hence the portrayal of the empirical domain as a subset of the actual domain would be incoherent if things were considered real but not actual.

However, Bhaskar also identifies causation with 'relatively enduring structures and mechanisms' that are "nothing other than the ways of acting of things" (Bhaskar, 1978, p. 14); or to put it in other words, "the generative mechanisms of nature exist as the causal powers of things" (Bhaskar, 1978, p.50) (Lawson, 1997, p. 21). Now for Bhaskar, causal powers and generative mechanisms exist in the domain of the real, but not in the actual, and this could be taken to imply that the same is true of entities. Since causal powers exist *only* as emergent features of entities, it is hard to see how these causal powers could exist in a different ontological domain from the entities of which they are features.

Fleetwood seems to imply something similar when he argues that "when… one writes that *a mechanism has a tendency to x*, one is, in reality, referring to the ensemble of structures, powers, and relations: it is, strictly speaking, the ensemble that has a tendency to *x*. Once understood, however, there is no harm in shortening the phrase by omitting reference to structures, powers and relations" (Fleetwood, 2001, p. 211). We can translate this into the language of emergence by equating 'ensembles' with higher-level entities whose components are lower-level entities and the relations between them. Fleetwood's argument thus translates into the claim that mechanisms are simply a level abstracted view of a multi-levelled entity. If a mechanism simply *is* an ensemble of structures, powers, and relations, then it *is* an entity – and it becomes yet clearer that if mechanisms are real but not actual, then so must be entities.

This, however, would seem to lead to a contradiction: one part of the argument entails that entities exist in the domain of the actual, whereas another seems to imply that they exist in the domain of the real but not in the domain of the actual. There are several possible responses to this contradiction; this paper will only discuss what seems to me the most plausible interpretation of Bhaskar's intention.[10]

This response is to insist that *actual* entities *do* possess *real* (but *non-actual*) causal powers – in other words, that a thing's 'way of acting' can exist in a different

[9] My thanks to an anonymous referee for reminding me of this point. Also see (Bhaskar, 1978, pp. 31-2).

[10] The version of this paper presented at the IACR 2004 conference offered a different resolution of this problem, which I no longer find plausible.

ontological domain from the thing itself. This would require a repudiation of Fleetwood's direct identification of mechanisms with ensembles, but could potentially be made consistent with Bhaskar's looser formulation. Given that the three domains of Bhaskar's ontology describe classes of what *exists*, this also rests on causal powers 'existing' in a somewhat different sense than entities, events, and experiences. I shall return critically to evaluate this way of resolving the contradiction in a later section, but first it is necessary to discuss in a little more detail the relation between the real and the actual in the critical realist account of cause.

Real causes and actual causation

Bhaskar's account of cause in *A Realist Theory of Science* is focussed on the role of causal mechanisms, which he identifies as part of the domain of the real (Bhaskar, 1978, p. 13). As we have seen, these mechanisms "exist as the causal powers of things" (Bhaskar, 1978, p.50). In such situations, we can, as Fleetwood suggests, work successfully with an abstracted ontology that ignores the fact that each entity or thing is composed of a variety of levels of lower entities, and simply sees it as existing at a specific level of organisation. The composition of the entities we seek to explain (or use as causal factors) is simply one of the many things that we abstract from in formulating laws. It therefore seems in the resulting generalisations that the entities which 'cause' and whose behaviour is 'caused' are free-floating level abstracted entities that are autonomous of their component parts, and that can be treated in those causal accounts as if they had no component parts at all.

Such a level abstracted conception of cause is perfectly usable and indeed positively desirable in the process of formulating theoretical laws. It also works quite reliably in many practical applications, both everyday and scientific, when level abstracted views of causation often seem to reflect what is going on well enough to provide us with reliable explanations and hence expectations. However, as Bhaskar himself recognises in more recent work, it is quite inappropriate for the discussion of what is happening over multiple levels when we turn to causation at the level of individual instances:

> "unlike theoretical explanation in at least many of the natural sciences, viz. from explanatory significant structures to their higher-order structural explanation, applied explanation of concrete singulars, like changes in a particular [entity], are a much messier affair … An event e at a level L is as likely to be (multiply) explained by elements at the same and lower-order levels in addition to higher-order (deeper) ones, and/or even laterally, diagonally, tangentially" (Bhaskar, 1993, p. 133).

Let me now use an example to show why level abstracted causal accounts are inadequate to the causal explanation of individual events over multiple levels. Consider the case of photosynthesis by a plant. In certain circumstances which need not detain us here, many plants 'convert' carbon dioxide from the atmosphere into oxygen. At the highest level of the event (i.e. a case of photosynthesis) we may simply say that it was caused by the power the plant has to photosynthesise. Many useful explanations may indeed rest on this power, and a scientist could investigate, for example, the differential rates at which plants produce oxygen in different contexts without worrying about how photosynthesis worked at the cellular or molecular level.

But there are some parts of the event concerned that would inevitably remain unexplained by such an account. At another level (the molecular), the process of

photosynthesis is a chemical reaction, and we could not explain either *how* photosynthesis works or *which* lower level parts of the entities involved are affected, and in what way, without looking at this process at the molecular level. This would not be an account of a different event, but a different account of the same event – one that is abstracted at a different level from the whole event.

And yet, the lower level account still gives us only a partial account of the causal process at work here, because any explanation at only the molecular level will miss the key *higher level* causal factors which are also necessary for the event to occur. Thus, these molecules would not have been brought into an arrangement that made this chemical reaction possible unless they had been organised into the form of the plant in the first place (with organisation into cells as an equally essential middle level). Furthermore, the same molecules blended into a soup would no longer have the causal power of photosynthesis, which arises from their organisation into the form of a plant. The causal power of photosynthesis thus belongs to the plant and not to the molecules, but to provide a complete causal explanation of what happens when photosynthesis occurs we need a causal account that operates at multiple levels simultaneously, invoking both the causal powers of the plant and the causal powers of its molecules.

In other words, it is impossible to explain fully the causation of the event except as the outcome of a causal interaction between the whole 'pyramids' – between the entities concerned, viewed in downwardly inclusive terms – and not just the single points at the top – the same entities viewed in level abstracted terms.[11]

We can see why this is a useful way to look at causation if we consider the problem posed to level abstracted accounts by multiple realizability, i.e. in cases where the higher-level outcome is consistent with a variety of different lower-level configurations. Such accounts are underdetermined, in that they can provide an account of the change that occurred at a higher level, but not an account of how the implicit lower-level changes occurred, thus leaving the higher level change floating unsecured without any confidence in how its components could have been brought to a state consistent with it. Downwardly inclusive accounts, by contrast, resolve this underdetermination since the whole range of states of all the component entities and sub-entities involved in the multi-levelled event are available to contribute to the causation of the lower-level changes.

Of course, each of the interactions at the lower levels can also be considered as inclusive events in their own right, so the higher-level event is at least partially composed of a whole set of smaller pyramidal events. Now as a result of this, a reductionist might argue that the inclusive account suffers from the opposite problem to that discussed in the previous paragraph: it may seem to be overdetermined,[12] if we believe that the higher level entities are no more than the sum of their parts, and lower level explanations are available for the behaviour of each of those parts. In this case, it would seem that causes at the higher level are redundant to the explanation of the event, since the lower level causes do all the causing that is needed to produce it.

[11] This is a sub-case of the determination of events in the actual by a mix of many causes; and also a case of what Bhaskar calls multiple determination, which will be discussed below.

[12] I use 'overdetermined' here, not in Althusser's sense, but rather to indicate a logically impossible case – i.e. where the set of causally effective factors exceeds those required to explain the set of outcomes, with the result that they appear to mandate a set of outcomes that may be inconsistent with each other.

But there are a number of problems with this reductionist claim. First, the meanings of some of the terms that describe the events to be explained may be incoherent at a lower level. Many animals can 'see' things, for example, but the concept of 'seeing' is meaningless when reduced to the behaviour of organs or cells. It is hard to see how lower level explanations of a concept that is meaningless at the lower level can be complete. Second, there will generally be features of the higher-level entity that are contingent on the relations between its components, and not just their separate presences summed; and it is the addition of these relations as an ongoing feature that distinguishes the higher level entity from the mere collection of lower level ones. A causal account of the lower level entities will not explain the higher level entity unless we go beyond that causal account to explain the set of relations between them too, and when we do this we have reintroduced the higher level entity into the explanation. Third, even the separate presences of those component entities in a particular situation is often difficult to account for except as the consequence of their being part of the higher level entity concerned. Why do we find this particular collection of lower level entities or events present in the first place, and not some other, perhaps random, collection? Why, because it is precisely this collection that constitutes the higher level entity and is held together by its morphostatic causes.

The causation of events thus operates across the whole pyramid of entities and sub-entities involved, not at a single level of it. Events, in all their multi-levelled glory, are the products of a combination of a variety of causal mechanisms operating on the prior state of the set of entities involved in the event. In Bhaskar's account, this 'individual instance causation' (which is of course interlinked with other individual instances of causation) occurs within the domain of the actual, but it is the consequence of the interaction of the real (but not actual) causal mechanisms or powers of the entities involved.

These causal powers exist as emergent properties of the entities that possess them. Because they emerge at a specific level (e.g. the ability to photosynthesis belongs only to the plant as a whole; the molecules or cells of the plant couldn't photosynthesise if they were not organised into the form of a plant), then it is entirely reasonable to think of them in level abstracted terms. Nevertheless, they can only lead to actual events when they are combined with a multiplicity of causal mechanisms from other levels of the ontological strata. Thus 'real' causal powers can be described in a level abstracted form, while 'actual' causation always occurs in the form of multi-levelled events. We may for some purposes be able to provide perfectly adequate explanations of these events that neglect many of the levels involved – perhaps even all but one. But when we wish to discuss questions concerning the relationship between different levels – such as the questions of emergence and reduction – we cannot do so in purely level-abstracted terms, but must recognise the *inherent* inter-relatedness of the different levels.

Multiple determination

Bhaskar himself addresses this question of the contribution of causes operating at different levels through a concept which he calls 'dual control', 'multiple control', or 'multiple determination'. In considering actual natural and social events, he argues, we must accept that different causal mechanisms and the interactions between them account for different aspects of the events concerned, and that no single law 'determines' the whole result:

> "Laws leave the field of the ordinary phenomena of life at least partially open... To say that laws situate limits but do not dictate what happens within them does not mean that it is not possible to completely explain what happens within them. The question 'how is constraint without determination possible' is equivalent to the question how 'can a thing, event or process be controlled by several different kinds of principle at once?' To completely account for an event would be to describe all the different principles involved in its generation. A complete explanation in this sense is clearly a limit concept. In an historical explanation of an event, for example, we are not normally interested in (or capable of giving an account of) its physical structure" (Bhaskar, 1978, pp. 110-111).

This is not just a statement about the relationship between different levels of stratification, but rather a more general discussion of the nature of actual causation. But Bhaskar makes the link to level stratification explicit in a more recent work: "Emergence makes possible the important phenomena of *dual* and *multiple control*" (Bhaskar, 1994, p. 75). It is precisely because 'the [actual] ordinary phenomena of the world' are inherently multi-layered, that we need to bring to bear different (real) single-layered causal mechanisms to explain different aspects of them. Thus explanation at each level, in the 'area of autonomy' left by the incomplete explanations at other levels, requires a 'putatively independent science' of that level (Bhaskar, 1978, p. 114). And it is in combining all these level-specific explanations of the different levels of a particular multi-layered event that we 'completely account for an event'. Although, of course, because we do not have viable sciences of every level, we can only produce incomplete subsets of the 'complete' multi-layered account, which is why such a complete account can be seen only as 'a limit concept'.

To put this in my terms: in decomposing the behaviour of a downwardly inclusive entity across its ontological levels, it is the organisation that appears at each level, the set of relations between the relevant lower-level entities, that is the 'extra' piece of explanatory information that appears at that level; and this is what makes the 'multiple determination' approach viable. We attribute a portion of the causal influence on a particular event to the set of relations between parts that constitutes the organisation of the topmost level, a portion to the set at the next level down, and so on. This allows us to construct causal accounts of multi-levelled single instance causation in which all the levels of the prior situation can have an appropriate influence on the various levels of the outcome. In this model, any insistence on 'explanatory priority' for any particular level becomes nothing more than a metaphysical prejudice.

It is worth noting that this conception of multiple determination is also required by the transcendental argument from the nature of experimental science. The most obvious causal regularity in experimental situations is the causal impact that the intervention of the experimenter has on the results of the experiment. Clearly there is a sense in which the experimenter 'causes' the results of the experiment (Bhaskar, 1978, p. 33). It is only when we have a concept like 'multiple determination' that allows different mechanisms at different levels to contribute to the determination of a multi-layered event that there is room for any other sort of cause to operate in experimental conditions as well as the causal input of the experimenter. Since experimental science works on the assumption that such other causes are in fact at work in experimental situations it also assumes that multiple determination is a feature of the world.

Multi-levelled causation of the actual, then, is an unavoidable feature of Bhaskar's ontology.

The consequences for Bhaskar's three domains

Let me now pull together and round out the implications of the foregoing for the three domains of Bhaskar's ontology.

There is relatively little to be said here about the empirical domain. In the context of level stratification, it is important to recognise that our experiences take the form of level abstracted views of a multi-levelled reality. This is the result of the combination of two factors: the inherent limitations of our perceptual tools, and the interpretive habits that are integrated into the very process of perception. Now, it is of course true that this form of perception is enormously effective in practical situations, or it would not have been favoured by biological evolution. And this effectiveness in turn derives from the fact that in many practical situations we can afford to ignore levels of stratification other than those we are in the habit of perceiving. Level abstracted perception and indeed level abstracted approaches to the causation of everyday behaviour work well for normal human purposes.

But science seeks to go beyond this type of understanding of the world we live in, and in delving into other layers of our level stratified world it reveals that there is more to events than meets the eye. An event at any given level is inseparably also made up of a set of events at lower levels (and may be a part of other events at higher levels). The theory of emergence enables us to see that if we want to explain a multi-levelled event then there will be a whole set of causal mechanisms involved, all operating simultaneously at multiple levels. If we wish to understand the relations between causes at different emergent levels we need to re-integrate these partial explanations with the other levels that are inseparably part of the same event, in what I have called a downwardly inclusive causal account. These multi-levelled events are the inhabitants of Bhaskar's domain of the actual. Other than qualifying the treatment of level stratification, then, the consideration of events here has no significant consequences for Bhaskar's three domains.

When we come to consider entities, however, the consequences are more significant. As we have seen above, there are good reasons to believe that entities belong in Bhaskar's domain of the actual – but there are also good reasons to believe that they belong in the non-actual portion of his domain of the real. Earlier, I introduced one possible approach to resolving this apparent contradiction, which seems in keeping with Bhaskar's intention. I will argue below, however, that this approach can not succeed.

The strategy adopted in this approach is to argue that entities belong unambiguously in the domain of the actual, but that actual entities can and do possess real causal powers. Here, the causal powers of an entity exist in a different domain from the entity itself. Since Bhaskar clearly locates the 'mechanisms' underlying causal powers in the non-actual real, he must mean something different than Fleetwood by the term mechanism. Rather than identifying 'mechanism' with the ensemble of parts and relations that constitute the (actual) entity possessing the powers, which would be incoherent if mechanisms are not in the actual, he must mean to identify 'mechanism' with something like 'the way in which the relations between the parts produce the causal powers of the whole'. This is indeed a useful description of the concept of mechanism, and exactly the sort of description we would expect in an emergent account of level-stratified reality.

Nevertheless it is clear that such a mechanism is not a separate thing from the entity that possesses it, but rather a consequence of how it is put together.

Furthermore, it is also clear that this conception of *real* mechanism is closely analogous to the process of *actual* causation.[13] In actual causation, any given event is the outcome of the actual interaction between the real causal powers of those entities causally involved in it, and the net outcome of these interactions depends upon the (purely temporary) relations in which these entities stand to each other at the time. This is directly analogous to the generation of the causal powers of a particular type of entity, which is the outcome of the interaction between the causal powers of its parts. The primary difference is that in the first case, the relations between the entities concerned are contingent and temporary, whereas in the second, the same set of significant relations is maintained over time as a result of the operation of morphostatic causes that maintain the structural stability of the entity, and hence there is a level of consistency in these causal powers over time.

Real and actual causation both therefore appear to be consequences of the same generic type of structural relation: the (diachronic) causal consequences that flow from a given set of entities existing (synchronically) in a given set of relations to each other. Real causal powers, on this account, are distinguished from ordinary instances of actual causation because (a) they provide only part of the explanation of any given event; (b) they may not be active in any given case; and (c) they are present consistently in all instances of the type of entity that possesses them.

The first and third of these differences, however, provide no obvious basis for declaring causal powers non-actual. In the first case, if the causation of an event occurs in the domain of the actual, we might reasonably expect the parts that combine to produce it to occur in the same domain. In the third, a consistent feature of an actual entity would seem at first sight to be actual too. It is the second difference upon which Bhaskar's argument rests: that causal powers constitute a separate ontological domain from actual causation because causal powers may be unexercised or unrealised in any particular actual case, and hence have an existence that is independent of actual *events*.

But once we recognise that entities are also part of the domain of the actual, as is done in the 'actual entities with real causal powers' argument under discussion, it is clear that the existence of a causal power can *not* be independent of the existence of an actual *entity*.[14] It is not the existence but only the *operation* of a causal power that is distinct from its instantiation in a particular actual entity. This suggests that the attribution of causal powers to a separate ontological domain overstates the separation between causal powers on the one hand, and the domain of actual entities, causation, and events on the other. Causal powers, like actual causation, are the consequence of the interaction of actual entities, and both are intimately tied to the existence of particular sets of entities.

The distinction between the two, however, is still important and useful. The identification of causal powers that can be combined in cases of actual causation becomes a technique for breaking down the analysis of causation into manageable chunks that can be separately investigated by scientists. This is indeed the obvious application of Bhaskar's analysis of cause, and I continue to believe that it is a very

[13] I develop the argument made here and in the remainder of this section in more detail in (Elder-Vass, 2005)

[14] Unless a 'real causal power' is an abstract universal. Although this is the approach I investigated in the version of this paper presented at the IACR conference, I now believe that this was not Bhaskar's intention, as is revealed by a careful reading of (Bhaskar, 1978, pp. 45-52). Bhaskar treats causal laws as universals, but not causal mechanisms.

valuable one, but it is hard to see how this justifies the treatment of causal powers as existing in a separate ontological domain from the actual entities that possess them or from the closely analogous process of actual causation.

My argument therefore suggests a methodological as opposed to an ontological division of actual causation from causal powers. This division still enables us to analyse cases of actual causation by identifying the entities involved and their characteristic emergent causal powers, then investigating how those powers combined to produce multi-levelled actual events. This may, after all, be all that Bhaskar hoped to achieve from the domain stratification of his ontology.[15]

Conclusion

I believe that many of the difficulties of existing approaches to emergence and reduction stem from the inappropriate application of a level abstracted ontology. I have sought to demonstrate that Bhaskar's depth ontology, as enhanced and, I hope, clarified in this paper, offers part of the solution to this problem. Conversely, I have also sought to show that the careful study of emergence can enrich Bhaskar's ontology.

On the one hand, depth ontology's division of actual events from emergent causal powers is an essential prerequisite for the understanding of causation in a multi-levelled world. On the other, when we need to take full account of the level stratification of the world, it seems that we need to see *actual* events and entities as inherently multi-levelled, whereas it is appropriated to see *real* causal powers (as well as *empirical* experiences) in level abstracted terms. It is thus only by combining causal mechanisms from a number of different levels that we can provide an adequate causal account of a downwardly inclusive event. This seems consistent with what Bhaskar anticipates in his account of multiple determination.

Having said this, it seems to me that there is still much more work to be done on this question of the relationships between and among levels in causal accounts, a question which is fundamental to understanding emergence and causation. Having understood that emergence is important, we need to examine in more detail just how it works, following in a tangential sense Bhaskar's own advice: "When a stratum of reality has been adequately described the next step consists in the discovery of the mechanisms responsible for behaviour at that level" (Bhaskar, 1978, p. 169). The outcome of such an exercise may well lead us to conclude that the higher stratum had not been adequately described in the first place. In a similar way, examination of the mechanisms of emergence may alter our perception of the nature of emergence itself and its ontological role, as I hope I have already demonstrated in this paper.

References

Archer, M. (1982). 'Morphogenesis versus Structuration'. *British Journal of Sociology*, 33, 455-83.

— (1995). *Realist Social Theory: The Morphogenetic Approach*. Cambridge: Cambridge UP.

[15] Although his argument is clearly also motivated by a desire to underpin his critique of positivism; I have not considered the implications of my proposed revision for that critique, but at first sight it would not seem to affect it.

Bhaskar, R. (1978). *A realist theory of science*. Hassocks: Harvester Press.
— (1993). *Dialectic : the pulse of freedom*. London ; New York: Verso.
— (1994). *Plato etc*. London: Verso.
Blitz, D. (1992). *Emergent evolution*. Dordrecht ; London: Kluwer Academic Publishers.
Buckley, W. (1998). *Society - a complex adaptive system*. Australia: Gordon and Breach.
Cilliers, P. (1998). *Complexity and Postmodernism*: Routledge.
Collier, A. (1994). *Critical Realism*. London: Verso.
— (1998). 'Stratified explanation and Marx's conception of history'. In M. Archer, R. Bhaskar, A. Collier, T. Lawson & A. Norrie (Eds.), *Critical Realism: Essential Readings* (pp. 258-281). London: Routledge.
Cunningham, B. (2001). 'The reemergence of 'emergence''. *Philosophy of Science*, 68, S62-S75.
Elder-Vass, D. (2005). 'Emergence and the Realist Account of Cause'. *Journal of Critical Realism*, 4, 315-38.
Emmeche, C., Koppe, S., & Stjernfelt, F. (1997). 'Explaining Emergence: Towards an Ontology of Levels'. *Journal for General Philosophy of Science*, 28, 83-119.
Fleetwood, S. (2001). 'Causal Laws, Functional Relations and Tendencies'. *Review of Political Economy*, 13, 201-220.
Lawson, T. (1997). *Economics and reality*. London ; New York: Routledge.
Lloyd Morgan, C. (1923). *Emergent evolution*. London: Williams and Norgate.
McLaughlin, B. P. (1992). 'The Rise and Fall of British Emergentism'. In A. Beckermann, H. Flohr & J. Kim (Eds.), *Emergence or Reduction?* (pp. 49-93). Berlin: de Gruyter.
Sawyer, R. K. (2001). 'Emergence in sociology'. *American Journal of Sociology*, 107, 551-585.
Sayer, A. (1992). *Method in Social Science*. London: Routledge.
Sperry, R. W. (1986). 'Macro- versus Micro-Determinism'. *Philosophy of Science*, 53, 265-270.