Social entities and the basis of their powers

Dave Elder-Vass
Department of Social Sciences
Loughborough University

Note: This is a pre-publication version of Elder-Vass, D. (2014) 'Social entities and the basis of their powers' in Julie Zahle and Finn Colin (eds) *Rethinking the Individualism-Holism Debate*, Springer, 39-53. It may differ slightly from the published version. Please refer to the published version for the final wording and pagination.

This paper offers an emergentist justification for the claims that certain sorts of social entities are causally significant, and that we should therefore reject methodological individualism, understood as a denial of the possibility of such structural causation.[1] [2] It argues that social structure is causally significant when there are social entities with emergent causal powers. Such powers are generated by processes of interaction between the characteristic set of parts of such an entity and the characteristic relations between them that occur in entities of the type concerned. The paper will offer both a justification of this argument and some brief illustrations. This argument, however, should also lead us to reject certain sorts of emergence theory, including much of the thinking about emergence that has been developed within the philosophy of mind. And it should lead us to question those forms of social structural discourse that *assume* structures are causally significant without *explaining* how the structures concerned can be significant in terms that are consistent with a coherent account of emergence.

Because the form of emergence theory advocated here is sometimes denigrated by philosophers on the grounds that it is too weak to ground causal claims, the paper begins by challenging the debate in the philosophy of mind that has conditioned them to expect something more from emergence theory than it can reasonably deliver. It then continues by outlining in general terms the *relational* form of emergence theory that I employ and making clear what can be claimed for it: that it provides a viable refutation of eliminative reductionism as a generalised strategy, where eliminative reductionism is defined as the denial of the causal significance of a category of entities or properties on the grounds that the causal work concerned is really done by some lower level category of entities or properties. In particular, I claim that this form of emergence theory can be employed to refute the form of eliminative reductionism that has come to be known as methodological individualism in the social sciences. The paper argues that this is a kind of emergence worth having, and indeed a kind of emergence that delivers just what we need from the concept: it justifies the

[1] I would like to thank Julie Zahle for her invitation to present this paper at the workshop 'Individualism, Holism, Explanation, and Emergence' in Copenhagen, November 2012, and the participants in that workshop for their stimulating questions and contributions.
[2] The paper draws heavily on arguments from my recent work (particularly Elder-Vass, 2010).

need for higher level sciences to study higher level mechanisms and powers, mechanisms whose explanation will never be made redundant by some lower level theory of everything.

This version of emergence theory supports a specific way of thinking about social structure that is arguably rather different from the ways that have tended to dominate sociological discourse. My project includes developing such theory; it is ultimately a sociological and not a philosophical project, and the final part of the paper illustrates how this way of thinking about emergence in the social sphere leads to what I take to be useful and interesting ways of reconceptualising social structure.

Emergence and the Philosophy of Mind

At the heart of my approach to emergence is the idea that there are things in the world (objects, or entities, if you like) that exert a causal influence and that emergence theory is about the relation between the causal significance of things and that of their material parts. At its simplest, this kind of emergence theory claims just that things have causal powers that their parts would not have if they were not organised into this kind of whole. Such emergent properties are thus similar to the concepts of *collective* or *systemic* properties.

Philosophers sometimes dismiss these as familiar and uninteresting, as if they were entirely unimportant by comparison with stronger forms of emergence. The pursuit of such forms is heavily focused on the philosophy of mind, and the search for a theory of emergence that will justify the claim that *mental properties* are special and in particular that they are in some sense causally autonomous of *physical properties*. I have gradually come to the conclusion that the entire debate on emergence in the philosophy of mind is shaped by a fundamental error that we may call *residual Cartesianism:* the belief that mental properties are radically different from other kinds of properties and thus in some sense exempt from the laws of causality.[3] Although the physicalist position within the debate is explicitly opposed to this belief, there is a sense in which it seems to be trapped within a terminology and a set of base assumptions that have already been distorted by an interlinked series of Cartesian moves.

The first of these moves is to abandon the notion of things or objects (which I shall call *entities*) and instead to develop the argument in terms of various kinds of *properties*. This is more or less a forced move if one is to defend residual Cartesianism, since it now generally agreed that there is no such *thing* as a mind (whereas there clearly is such a thing as the brain), and so the mental and the concept of mind are rescued by using these terms to refer to a set of *mental properties*. Connected to this move, the relations between entities are also relabelled as just another variety of properties.

The second move is to separate mental properties from physical properties, a move that Searle has labelled *property dualism* (Searle, 2002), and again a necessary move if one is to preserve the idea that mental properties could be fundamentally different from others. This notion of a physical property is however problematic. There is something rather odd, for example, about the fudge of saying that physical properties are properties studied by physics and then immediately going on to treat neurological properties – which generally are not studied by physics – as physical properties. Taking a different approach, we can usefully talk about physical *entities*, in the sense of material

[3] The term *residual Cartesianism* has been used elsewhere in the literature, although generally with a different residue in mind (e.g. Allen et al., 2012).

entities, but it would be equally odd to talk about *properties* being physical in the sense of material. What is material about being green, for example? In practice, then, the concept of a *physical property* in this debate simply means "all properties that are not mental properties": it is a fake other that serves only to legitimate the idea that mental properties are *different* without providing any substantive specification of what it is that they are taken to be different from. Indeed, there seems to be nothing that so-called physical properties have in common that they do not also share with mental properties.

But it would be laughable today to insist that mental properties are entirely unconnected to physical entities in the form of brains, and so residual Cartesianism must make a third move that somehow recognises this relation while nevertheless asserting the independence of mental properties. This is achieved by treating the physical as a *lower level* with an ambiguous relationship to *higher level* mental properties: one in which the nature of these levels is left opaque, and some mixture of dependence and autonomy between them is asserted. This set of moves is not only assumed by emergentists, but also by many of those who question the emergence of the mental. It is built into both the programme and the language of a significant proportion of the philosophy of mind – or at least of that perhaps rather limited section of the philosophy of mind with which I have engaged.

I suggest that all three of these moves are errors. The properties that are of interest in debates on emergence are always properties *of* specific entities or types of entity, and it matters to these debates *which* entities they are properties of. These features are obscured by ignoring entities. Similarly, the relations between entities are quite a different class of phenomenon than properties that are possessed by a single entity in isolation from others. Again, this is a distinction that matters to this debate, as we will see, and it is one that is obscured by the exclusive focus on properties. Without distinguishing between entities, their properties, and the relations between them, I will suggest, we cannot even make sense of the problem of reductionism, let alone possible solutions to it.

It is not unreasonable to suggest that there are phenomena it is useful to label as mental properties: beliefs and desires, for example. But, like Searle, I am baffled by the idea that these are somehow distinct from so-called physical properties (Searle, 1992: chapter 1; Searle, 2002). For me, mental properties are a subset of the properties of material entities in much the same way that colour properties or electro-magnetic properties are a subset of the properties of material entities. Once we accept this, there is no need to abandon the language of entities and their properties, and no need to struggle with the underpinnings of the separation between mental and physical properties. A mental property is a property of a material entity, a person, just as the capacity to walk or sing or a certain skin colour may be a property of a person. None of these properties is any more or less *physical* than any of the others, and there is therefore no reason for us to find the relation between mental and physical properties as such problematic.

Nor is there any reason to think that mental properties are higher level than other ('physical') properties in general. The concept of levels is meaningful within an entity-based ontology, since there is a compositional hierarchy within any given entity, and it is useful to think of wholes as being at a higher compositional level than their immediate parts, those parts being at a higher level than their own parts, and so on. Given such an ontology, we may call the properties of a given whole *higher level* properties than the properties of its parts. But without such an ontology, it is not at all

clear what it would *mean* to call a set of properties *higher level*. Given an entity-oriented ontology, we could say that mental properties, being properties of whole human beings, are higher-level properties than the properties of, for example, neurons, but *not* that they are higher-level properties than other ('physical') properties of whole human beings.

There does, nevertheless, remain a question to which some form of emergence theory might be relevant: What is the relation between a person's beliefs, desires, etc., and that person's neurological structure? Unlike the debate over mental and physical properties, as it has usually been conceived in the philosophy of mind, this is much the same sort of question as many others we may encounter in scientific contexts, such as "What is the relation between the radiation that may be generated by a star and its physical structure?" or "What is the relation between the capacities of an organisation and its structure?" In some respects, all of these questions are scientific, as opposed to philosophical, questions. I am one of those optimists who believe that eventually, if we don't destroy ourselves first, humanity will develop satisfying, useful, and largely accurate scientific explanations of all of these questions. What, then, might be the relevance of emergence theory?

It is useful, primarily, as a response to eliminative reductionists: those who question a whole set of explanations on the grounds that they appeal to the causal significance of a category of entities that they believe to be causally epiphenomenal, usually because they believe that the causal work is really done by some lower level category of entities. This is arguably a problem of the immature sciences. Practitioners of the mature sciences rarely need to worry about eliminative reductionism since in these sciences it is clearly useful to proceed as if entities at a variety of compositional levels have causal significance and there is no obvious reason to doubt that this is so. But sociologists need to worry about it because it is still considered plausible by some to argue that social structures are not causally significant and therefore that we should reject any causal explanation that invokes them. And perhaps psychologists need to worry about it in an age that threatens to become neurologically reductionist. It is therefore useful to have a general account of emergence that may be used in response to such reductionisms and applied to specific cases.

Relational emergence theory

As an introduction to such an account, consider what we might need by way of a response to eliminative reductionism. Once we discard residual Cartesianism, there is no need to seek to sustain its continued yearning for realms of existence that are in some sense exempt from causal explanation and yet at the same time capable of exerting a causal influence. We can therefore reject another deeply ingrained philosophical presumption, one that is wrapped up in many uses of the concept of reductionism. This is the assumption that if we can offer a 'lower level' explanation for something, then that something cannot itself have causal significance.

One argument that is sometimes deployed in this context is Kim's causal exclusion argument (e.g. Kim, 1993: 203-9; Kim, 2006: 558). Kim argues that if there is a mental property and a physical property that realises it, and if we accept that the original physical property causes a change to a new physical state, then we cannot also believe that the original mental property causes a change to the corresponding new mental state. In some versions of the argument, he makes clear that the physical property he has in mind is what he calls a *microstructural* property, which corresponds to a "unique complete microstructural description: that is, any physical system can be exhaustively described in terms of (i) the basic particles that constitute it… ; (ii) all the intrinsic properties of these

particles; and (iii) the relations that configure these particles into a structure" (Kim, 1999: 6). A microstructural property, in other words, is a translation into property language of the concept of an entity: the microstructural property of a system is a complete description of all the parts and relations that constitute it as a system, and so in specifying a microstructural property we are essentially describing the entity concerned. My summary of his argument, then, could be rephrased as follows: if there is a mental property that is a property of a material person in a given state, and if we accept that the original state of the material person causes a change to a new state of the material person, then we cannot also believe that the original mental property causes a change to the corresponding new mental state.[4]

Whatever this argument implies, it does *not* imply that higher level properties, in the compositional sense of *levels* advocated above, are not causally effective if lower level properties are. As Kim points out, microstructural properties are not lower level properties but macro properties: "macro since it belongs to the system as a whole constituted by the system's basic micro-constituents, their intrinsic properties, and the relations that structure them into a system with unity and stability as a substance" (Kim, 1999: 7). Only the whole system, the whole entity in my terms, can have the microstructural property, so it is not a lower-level property in the compositional sense of *levels*. His causal exclusion argument is not about the relationship between the causal power of entities at different compositional levels, but about the relationship between the causal power of an entity at a given level and a property at the same level. Indeed Kim himself appears to advocate a view that is similar to the account of emergence offered in the current paper, though he does not *call* it emergence: "the fact that we can micro-structurally explain why a micro-based property has a certain set of causal powers does not mean that these causal powers are identical with the causal powers of its micro-constituents. Micro-reductively explainable causal powers may be new causal powers, net additions to the causal structure of the world" (Kim, 1998: 117). Kim's argument therefore cannot be invoked against relational forms of emergence theory such as that advocated here.

Let us turn, then, to the ontological basis of relational emergence theory. My base assumption is that our universe is populated with stuff that interacts and as a result forms progressively more complex structures, which sometimes have a degree of stability and persistence, and which in turn interact in progressively more complex ways. It is populated, as a result, by structured entities that we can think of as decomposed hierarchically into parts and sub parts at various levels. Those complex structures can have effects that the same stuff cannot and does not have when it is not organised into such structures. In the sense of emergence advocated here, this means that entities at each level can have emergent causal powers: powers to affect the world that would not be possessed by the parts if they were not organised into entities of this nature.

It is the interactions of the part-stuff that produce the effects of these structures but those interactions depend on the set of relations between the parts that only exist when this structure is present, in other words when the set of relations that is required to organise these parts into this kind of structure is operative. For this reason, I refer to this as *relational* emergence theory. In a sense, it is the relations between the parts of any given whole that provide a bridge or intermediate

[4] The reader may substitute 'brain' for 'material person' if this makes the argument more accessible. But mental states may depend on our entire nervous system and not just our brains.

level between one compositional level of entities and the next: this is a further input into the properties of the whole that is not in itself explained when we explain the intrinsic (that is non-relational) properties of the parts.

In principle, science may provide explanations of how these interactions produce the emergent causal powers concerned; that is, it may identify the *mechanisms* that generate the higher level powers. Such explanations, however, are not reductions in what I take to be the core (that is, eliminative) sense of the term since they do not entail that the higher level structure is not required for the production of the effect. At the risk of confusing things by invoking a term that is often thought to have eliminative connotations, we could call these explanations *explanatory reductions*, but they do not entail the eliminative reductionist conclusion that the whole entity is somehow causally irrelevant.

A common response to such arguments by those of an eliminative reductionist frame of mind is that in such cases it is 'really' the parts that are doing the causal work. Let me call this position *reallyism*. The incoherence of reallyism is easily exposed by applying it recursively. If, for example, it were correct to argue that it isn't 'really' an organisation that is exerting a causal influence, but the organisation's members, then the same conceptual frame could be applied to the members themselves and would lead to a series of further conclusions: that it isn't really the members but their cells that are doing the causal work, not really the cells but their molecules, not really the molecules but their atoms, and so on until all causal power drained away into the bottomless pit of our scientific understanding of the most fundamental structures underlying the objects of particle physics (if indeed there are such structures) (cf Block, 2003). To sustain the kind of claim advanced by many eliminative reductionists – that *some* structured entities have causal powers and others do not – we would need a further ontological analysis that provides a justification for discriminating between forms of structure that may and may not have such powers. It is rare to find any such argument from reallyists, or indeed any recognition that their position requires it. One exception is the social psychologist Rom Harré, who argues that humans occupy "two realities… biology and conversation", characterised by causal and symbolic relations respectively (Varela and Harre, 1996: 317), and that in each reality all explanations can be reductively expressed in terms of one and only one kind of "powerful particular". In the former space, the powerful particulars are electrons (or fields), and in the latter they are persons (Harré and Bhaskar, 2001: 31). While I find Harré's account unsatisfactory, he does at least strive to be consistent with the internal logic of reductionist arguments. More typically, reallyists appear to base the belief that certain sorts of entities (generally people) must be causally significant, despite the logical structure of their own argument, on little more than an anthropocentric prejudice.

My general ontological argument, then is that entities have emergent causal powers when they are capable of exerting an influence on the world that their parts would not be able to exert were they not organised into such a whole.[5] Such powers are produced by mechanisms, processes in which the parts of the entity interact to generate the influence concerned, and we may be able to explain such mechanisms scientifically. This, however, does not alter the fact that these powers would not exist if the whole did not exist, and therefore we may conclude that these are causal powers of the whole and not of the parts. Without such a move, it is difficult to see how we could locate causal power

[5] This argument draws significantly on the early work of Roy Bhaskar (Bhaskar, 1975; Elder-Vass, 2005).

anywhere in the world; with it, we have a coherent way of examining both how causal power is generated and how powers interact to produce events.

In terms of the philosophical literature, then, I am advocating what would usually be labelled a rather weak form of emergence. Stephan, for example, distinguishes between *strong* forms of emergence theory, such as Broad's, which assert that emergent properties cannot be explained in terms of lower level parts and the relations between them, and *weak* forms, which do not make this claim (Stephan, 2002: 79).[6] It is strong forms that are sought by residual Cartesians; whereas in Stephan's terms relational emergence is a variety of the weak form. Stephan, however, goes on to argue that weak forms are "compatible with reductionistic approaches without further ado" (Stephan, 2002: 79). Such arguments must be carefully qualified by reference to the *kind* of reductionism at issue. I am tempted to suggest that we might also usefully distinguish between *weak* and *strong* forms of *reductionism*: weak reductions are simply explanations of a property of an entity, whereas strong reductions are explanations of such a property made entirely in terms that are compatible with the non-existence of the whole structured entity. Weak reductions do not entail that the property being explained or the entity possessing the property can be eliminated from a viable explanation of the effects of the property; whereas strong reductions do: they are eliminative reductions in the sense defined earlier. Relational emergence theory, we may then say, is thoroughly compatible with weak reductions of emergent properties but not with strong reductions of emergent properties (although it is compatible with strong reductions of non-emergent, i.e. resultant or aggregative, properties) (Wimsatt, 2000).

Philosophers of mind are prone to question the value of emergence theories that are compatible with weak or explanatory reductions, since they appear to be in search of an emergence theory that will exempt mental properties from 'physical' explanations. But once we discard residual Cartesianism we have no need for strong theories of emergence. What we *do* need is ontological theories that are compatible with the successful practices of actual science, and thus with us living (as we do) in the kind of world in which such practices *can* be successful.[7] In particular, we need theories that are compatible with two key features of that practice. First, a vast range of entities is treated as having causal significance; and second, this causal significance is taken to be compatible with the production of explanations of how it arises (Gell-Mann, 1995: 112). The kind of emergence theory advocated here meets both of these requirements, and provides us with an ontological framework that recognises the need for sciences of each level of structure: sciences that recognise which macro structures have which kinds of causal influence and also seek to explain how they can have it. This is all we need from the concept of emergence. And this gives us all that we need by way of resisting eliminative reductionism: it enables us to justify the assertion that higher level entities have causal powers while resisting the anti-scientific insistence that such powers are in some sense uncaused or unexplainable. This kind of emergence theory may be of no use to those seeking to substantiate pseudo-Cartesian dualisms in defence of concepts of mind that are somehow to be exempted from the normal processes of causality and explanation. But however it is labelled, it is *strong enough* to support a coherent approach to causal explanation.

[6] Other writers have given somewhat different meanings to these terms (e.g. Bedau, 1997).
[7] I take this to be one of the core arguments of Bhaskar's *Realist Theory of Science* (1975).

Causal powers of social entities

My interest in emergence theory, ultimately, is not philosophical. The focus of my work lies in the area where the ontology of the social world meets empirically relevant sociological theory. For me, then, emergence theory is a tool for making sense of concepts like social structure and agency, in ways that can be productive for sociological theory. Does it, for example, provide us with a coherent response to approaches like methodological individualism, a response that works not just as a philosophical argument, but because we can use it to develop accounts of social structure that show *how* it can be causally significant? Developing such accounts has been the primary focus of my work so far, and I hope to show that this does indeed offer at least part of a refutation of methodological individualism. But it also has disruptive implications for previous understandings of social structure, which are at least as important.

Perhaps the simplest way to apply emergence theory to these issues would be to assert that social structures, as they are already understood in existing sociological work, are, or are produced by, emergent properties. Such an approach would bolt on an emergentist justification to an existing body of thinking about social structure in a superficially non-disruptive way, but it would also generate a range of challenging questions. The most striking of these arises from the claim that emergent properties, as I develop the argument, are properties of entities that arise from the interactions between their parts. If this is so, then a coherent bolting on of emergence theory would require us to identify what entities, parts and interactions are involved in existing conceptions of social structure. While there are many ways in which one could attempt to fit these concepts together (Elder-Vass, 2010: chapter 4) perhaps the mainstream assumption has been that the entities involved in social structure are *societies*. We might then regard social structures as emergent causal powers of these societies. But there are a number of reasons why such an approach is unsatisfactory.

The most obvious is that the concept of a society is itself rather contentious. It carries with it a series of assumptions, beginning with the assumption that territorial states are taken to be boundary-defining for societies – an assumption that may be labelled *methodological nationalism* (Chernilo, 2007). But if this is the case, and if it is societies that are the causal agents at work when social structure is invoked as an explanation, this would seem to imply that social forces should operate similarly throughout the territory of any given state, and differently in the territories of other states. This is problematic because the geographical scope of specific social structural forces clearly varies enormously. Many are influential in spaces that are much larger than the territory of particular states, such as the normative standards espoused by various world religions, or the influence of Microsoft on practices of interaction between humans and computers. Many others are influential in much smaller spaces, such as separatist political parties and very local cultural traditions such as some language dialects. And still more are influential in spaces that are simultaneously both larger in some dimensions and smaller in others than the territory of particular states, such as the New Age or hip-hop subcultures. If social structures were an emergent property of whole societies, such variations in scope of influence would be incomprehensible, as would be the widespread phenomenon of conflicting social structures within the territory of the same state – employers organisations and trade unions, or meat-eating and vegetarianism, for example.

It is far more plausible to argue that the entities at work when we talk of the causal influence of social structure are not societies at all, but smaller (in most cases) and more specialised social

entities. Rather than a few monolithic 'societies' exercising social causative powers, what we find is a vast range of what we may call meso-level social entities (Stones, 2005). Perhaps the most obvious type of such a social entity is organisations, and we may usefully consider how the causal model advocated here alters theoretical understandings of organisations (Elder-Vass, 2010: chapter 7).[8] An organisation is a social entity, whose parts are primarily people, in a structured set of relationships to each other. Those relationships are often characterised as the roles of the individuals concerned. As a result of its members being committed to interact in the ways specified in their roles, the organisation has the capacity to have a causal impact on the world that its members would not have if they were not parts of the organisation concerned. Thus, for example, an orchestra has the causal power to produce harmonious music, a power that is generated as a result of the musicians (and arguably their instruments) who are its parts interacting in the ways specified in their roles (violin player, pianist, conductor, etc). If they were not organised into an orchestra, the players and instruments would not have the causal power to produce harmonious music.

However, the effects of the relational model of emergence are rather more radical if we consider a different class of social entities: the group of entities I have called *norm circles*. Norm circles, I argue, are the type of social entity that is causally responsible for normative social institutions or social practices. A social practice may be defined as a recognisable pattern of behaviour or action that occurs repeatedly in a social space, ranging from something as simple as standardised forms of greeting to something as complex as the practices that surround (and indeed create) what we think of as property. Not only simple day to day practices like queuing, but also phenomena like money, religion, language, culture and indeed organisations depend on normatively standardised practices.

Sociologists often invoke the concept of a normative social institution as a kind of explanation of the regularity of such practices. What is generally accepted is that such institutions are driven by normative pressures (socialisation, for example) which encourage people to conform to the practice and may penalise those who do not. What is not generally accepted, and indeed has been the focus of over a century of debate, is just what form social institutions take that gives them the capacity to exercise such an influence.

The hypothesis examined in my work is that social practices are produced primarily by the causal power of social entities that I call norm circles (Elder-Vass, 2010: chapter 6). A norm circle is the group of people that are committed to endorsing and enforcing a specific norm, a specific standard of observable behaviour. The relation between them that gives them the collective capability to influence behaviour – a greater influence than an unconnected group of individuals would have – is the sense of shared commitment they have to supporting the norm. The members of a norm circle may be unaware of the full extent of the group, and they may not even think of it as a group, but they are generally aware when they act in support of a norm that they are not simply expressing a purely idiosyncratic personal attachment to a particular standard of behaviour. Rather, they are aware that when they do so they are endorsing a standard that others also endorse, and often do so with the expectation that others would support and approve of their action. The individual, in other words, has a sense, however vague and minimal, that she is acting on behalf of something wider than herself when she acts in support of a norm, and that sense increases the likelihood that she will

[8] The remainder of this paragraph and the following six are drawn largely from Elder-Vass (2012b).

act in its support, by comparison with the isolated individual with a purely personal attachment to the standard of behaviour concerned.

This sense, in turn, is a product of the same process that tends to encourage conformity to the norm – the generative mechanism that underpins the power of a norm circle to increase such conformity. The heart of this process is repeated exposure of individuals to acts of endorsement and enforcement of the norm concerned. If, for example, I repeatedly see people criticising those who try to jump queues, I will start to understand the norm of queuing, and to believe that I face an environment in which I will be sanctioned negatively if I fail to observe it. I will, in other words, develop beliefs about my normative environment which will tend to lead me to conform to the norm of queuing in the future, as a result of the actions of members of the norm circle for queuing. Here, then, social structure – the norm circles that produce the normative environment – is exerting a top-down influence on individual action.[9] Similar effects can be produced without us forming conscious beliefs; most of us, for example, understand and implement the norms prevalent in our social space regarding how close one should stand to someone when talking to them, even though these norms are rarely stated explicitly and are mostly endorsed and enforced rather subtly by non-verbal signals.

Norm circles, then, operate through individuals. On the one hand, it is the actions of individual members of the norm circle acting in support of a norm that signal the normative environment to other individuals and these individual actions, therefore, reproduce and/or transform these social structures.[10] On the other, those pressures do not lead directly and mechanically to norm conformity but rather influence the stored beliefs and dispositions of the affected individuals, which then in turn influence their subsequent behaviour. Nevertheless, I argue, the resulting increase in the tendency of those affected individuals to conform to the norm is causally influenced by the norm circle, and not just the individuals. We may compare this with the case of a person who switches on a torch by pressing the switch with her finger: although her finger presses the switch, it could not do so except as a result of the causal influence of the whole individual, and so the event of switching on the torch is caused by a power of the whole individual and not just of the finger. In a similar way, the norm circle can only influence us through its individual members, but those individual members would not influence us in that way, or at least not as strongly and as often as they do, if they were not part of a wider norm circle, and so their act of influencing us is produced by a causal power of the norm circle and not just of the individual. Just because the causal power of the norm circle is exercised through individuals, this does not mean that it is *really* a causal power of individuals, since it is a causal power that would not exist if those individuals were not organised into a norm circle.

There is much more to be said about the theory of norm circles. For example, norm circles in contemporary societies are diversely intersectional – different norms are supported by different but profusely overlapping groups of people. This in turn makes normative change more likely than in more homogeneous societies, since individuals are open to the influence of competing norm circles and may move between them. One implication is that the theory of norm circles is not merely concerned with the reproduction of a stable normative environment. It seeks to explain how normative influences contribute to the production of social actions that conform with prevailing practices, and thus may contribute to the reproduction of the normative environment, but there are

[9] This is therefore equivalent to the first half of Margaret Archer's morphogenetic cycle of interaction between structure and agency (Archer, 1995).
[10] As in the second half of Archer's morphogenetic cycle.

many reasons why norms may be transformed rather than reproduced in some social situations. Another feature of this approach is that it is compatible both with cases in which individuals conform with and/or endorse norms as a result of internalising a strongly moral sense that they are right, and also with cases in which they do so for much more instrumental reasons, such as seeking approval and avoiding punishment.[11] Such arguments may also be extended to a variety of other normatively shaped phenomena, such as discourse, language, and knowledge (Elder-Vass, 2011; Elder-Vass, 2012a).

For the purposes of this paper, however, the relevance of norm circle theory is that it illustrates how the conception of social entities with emergent causal powers may be deployed to explain causal influences that have previously been attributed to rather amorphous conceptions of social structure, and to explain them in a rather less mystifying fashion than these earlier theories of structure. This has a number of important implications for the questions at issue here. First, it substantiates the notion that there may be social entities that have emergent causal powers, powers that would not exist in the absence of the entities concerned, and therefore casts doubt on methodological individualisms which argue that the idea of structural power is really just a misdescription of the powers of individuals. Second, it provides a clear explanation of how such emergent causal powers may come about as the result of interactions between lower level parts of the social entities concerned, including (but not necessarily restricted to) human individuals. It thus provides a weak reduction of such powers while also providing grounds for denying that form of strong or eliminative reductionism we know as methodological individualism, since it describes a form of social causal power that depends on the existence of specific sorts of groups, and not just on the existence of human individuals.

Conclusion

This paper has argued that conceptions of emergence derived from the residual Cartesianism that has shaped the debate in the philosophy of mind have misled philosophers about what can usefully be achieved with emergence theory. This has led them to set a standard of unexplainability for emergence theories that is both unachievable and also inconsistent with our scientific understandings of the world. A more realistic and achievable objective for emergence theory, and one that is more consistent with scientific practice, is that it provides a justification for ascribing causal powers to entities while allowing that we may also be able to develop causal explanations of such powers. The form of relational emergence theory summarised in this paper does provide just such a justification and this is all that we need from emergence theory.

This is also, however, a form of emergence theory that has significant implications for the social sciences. It provides a clear justification for rejecting methodological individualism. But it also prompts an approach to social explanation in which we identify the social entities that are exercising causal powers, and the mechanisms through which they do so. Such an approach has the potential to bring explanatory clarity to a number of areas where it has been lacking in the past; this paper's discussion of norm circles as the social entities behind normative social institutions illustrates the point. It is an approach, therefore, that does not simply validate existing theories of social structure, but rather requires a new approach to theorising social structure, one that includes examination of the processes of interaction between individuals that generate social powers.

[11] For further discussion of these and other complexities see Elder-Vass (2010: chapter 6).

Bibliography

Allen, C., Todd, P.M., & Weinberg, J.M. (2012). 'Reasoning and Rationality'. In E. Margolis, R. Samuels & S. P. Stich (Eds.), *The Oxford Handbook of Philosophy of Cognitive Science* (pp. 41-59). Oxford: Oxford UP.

Archer, M.S. (1995). *Realist Social Theory: The Morphogenetic Approach*. Cambridge: Cambridge UP.

Bedau, M.A. (1997). 'Weak Emergence'. *Philosophical Perspectives*, 11, 375-399.

Bhaskar, R. (1975). *A realist theory of science*. Leeds: Leeds Books.

Block, N. (2003). 'Do Causal Powers Drain Away?' *Philosophy and Phenomenological Research*, LXVII, 133-150.

Chernilo, D. (2007). *A social theory of the nation-state*. Abingdon: Routledge.

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

Elder-Vass, D. (2010). *The Causal Power of Social Structures*. Cambridge: Cambridge UP.

Elder-Vass, D. (2011). 'The causal power of discourse'. *Journal for the Theory of Social Behaviour*, 41, 143-160.

Elder-Vass, D. (2012a). *The Reality of Social Construction*. Cambridge: Cambridge UP.

Elder-Vass, D. (2012b). 'Top-down causation and social structures'. *Interface Focus*, 82-90.

Gell-Mann, M. (1995). *The quark and the jaguar*. London: Abacus.

Harré, R., & Bhaskar, R. (2001). 'How to Change Reality: Story v. Structure - A Debate'. In J. López & G. Potter (Eds.), *After postmodernism : an introduction to critical realism* (pp. 22-39). London: Athlone Press.

Kim, J. (1993). 'The Non-Reductivist's Troubles with Mental Causation'. In J. Heil & A. R. Mele (Eds.), *Mental causation* (pp. 189-210). Oxford: Clarendon 1993.

Kim, J. (1998). *Mind in a physical world*. Cambridge, MA: The MIT Press.

Kim, J. (1999). 'Making Sense of Emergence'. *Philosophical Studies*, 95, 3-36.

Kim, J. (2006). 'Emergence: Core ideas and issues'. *Synthese*, 547-559.

Searle, J.R. (1992). *The rediscovery of the mind*. Cambridge, MA: MIT Press.

Searle, J.R. (2002). 'Why I am not a Property Dualist'. *Journal of Consciousness Studies*, 9, 57-64.

Stephan, A. (2002). 'Emergentism, Irreducibility, and Downward Causation'. *Grazer Philosophische Studien*, 65, 77-94.

Stones, R. (2005). *Structuration theory*. Basingstoke: Palgrave Macmillan.

Varela, C.R., & Harre, R. (1996). 'Conflicting Varieties of Realism'. *Journal for the Theory of Social Behaviour*, 26, 313-25.

Wimsatt, W.C. (2000). 'Emergence as Non-Aggregativity and the Biases of Reductionisms'. *Foundations of Science*, 5, 269-297.