

## Article

## Open Access

Dave Elder-Vass

# Collective Intentionality and Causal Powers

DOI 10.1515/jso-2014-0039

**Abstract:** Bridging two traditions of social ontology, this paper examines the possibility that the concept of collective intentionality can help to explain the mechanisms underpinning the causal powers of some social entities. In particular, I argue that a minimal form of collective intentionality is part of the mechanism underpinning the causal power of norm circles: the social entities causally responsible for social norms. There are, however, many different forms of social entity with causal power, and the relationship of collective intentionality to these causal powers varies, depending on the form of the mechanism underpinning the power concerned. Some powers depend on collective intentionality, and others do not.

**Keywords:** Collective intentionality; Institutional reality; Critical realism; Norm circles; Searle.

## 1 Introduction

This paper examines the possibility that collective intentionality might play a role in the existence and powers of social entities. In doing so, it brings together two paradigms or traditions of work on social ontology, centred on the concepts of collective intentionality and causal powers respectively, in the hope that a productive synergy might be found between them. The collective intentionality tradition, notably including John Searle, Margaret Gilbert, Raimo Tuomela, and Michael Bratman, focuses on whether and how there could be collective intentions equivalent in some way to the intentions of individuals, and how this might contribute to the ontology of social institutions. The causal powers tradition is represented most strongly in the field of social ontology by critical realist thinkers, notably Roy Bhaskar, Margaret Archer, and Tony Lawson. For critical realists, causality rather than intentionality is the primary focus of ontological attention,

---

Dr. Dave Elder-Vass, Department of Social Sciences, Loughborough University, Loughborough, Leicestershire, LE11 3TU, UK, e-mail: d.elder-vass@lboro.ac.uk

 ©2015, Dave Elder-Vass, published by De Gruyter.

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 3.0 License.

and the concept of emergence is invoked to justify attributing causal powers to social structures. But how might these two traditions relate to each other? Are they competing paradigms, or complementary schools of thought, and if they are complementary, how might they be brought together productively?

A number of critical realist writers have already argued that these two perspectives are complementary (Smith 2010, p. 210–212; Elder-Vass 2012, chapter 4; Lawson 2012). Building on these contributions, and in particular on Elder-Vass (2012), this paper investigates one set of questions that arise when we seek to combine them. Given that I write from a critical realist perspective, the orienting concern of the paper is the issue of causal power, and in particular it focuses on the question of what role collective intentionality might play in the mechanisms that underpin the powers of social entities. I take it that this is not a question with a single simple answer, on the grounds that there may be multiple forms of collective intentionality, multiple forms of social entities, and thus multiple possible relations between these varying forms. This paper sets aside yet further complexities by conducting the argument in terms of social entities rather than social structures, thus evading consideration of the relation between the many and varying treatments of social structure in both philosophy and social science, which I have discussed elsewhere (Elder-Vass 2010). As a secondary concern, the paper touches briefly on the ways in which Searle's version of collective intentionality might itself require support from a causal powers analysis. It thus points towards the possibility that *both* traditions might benefit from their articulation.

The paper begins by introducing collective intentionality through a brief discussion of the work of John Searle and then summarises the critical realist perspective on social ontology. With this background in place, it illustrates the central question of the paper by considering the role of rather a minimal form of collective intentionality in the mechanisms underlying social norms. The paper then discusses a range of other possible relations between collective intentionality and social causal power by briefly considering some other types of social interaction. In doing so, it seeks to make stronger connections between the concept of collective intentionality and a variety of empirical social phenomena (thus also linking philosophical social ontology with issues in sociological theory) though it still operates largely at the level of abstract types rather than specific empirical cases.

If the argument of this paper is sound, there is significant value to be had by bringing together these two traditions, though doing so may demand developments in both. These include the need to recognise multiple forms of collective intentionality with differing roles in different types of social entity, and the need to consider whether and how collective intentionality contributes to the various mechanisms that give these social entities their causal powers.

## 2 Searle and Collective Intentionality

Although there are many scholars writing on collective intentionality, such as Gilbert and Tuomela who did important work before him and indeed continue to do so, it is the work of Searle that is probably most prominent, particularly outside the field of philosophy of social science. Even Searle's work has limited points of contact with social science itself, which he has tended to ignore in his writing (Osborne 1997, p. 98; Hund 1998, p. 130; Wettersten 1998, p. 132). Nevertheless, it provides the most widely accessible version of collective intentionality theory and thus a convenient place to begin our discussion of this tradition.

As Searle's position is well known I will be brief about the details. The focus of his argument in *The Construction of Social Reality* (Searle 1995) and *Making the Social World* (Searle 2010) is the concept of *institutional reality*, and in these books he develops what we may call a moderate constructionist account of institutional reality. Searle argues that there are brute facts of the world that exist independently of social interaction, but on top of these we construct institutional facts, and institutional reality consists of this set of institutional facts. In general, he argues, an institutional fact is created when a *constitutive rule*, supported by *collective intentionality*, assigns a new *status function* to a previously existing fact. This general principle explains the existence and nature of a vast range of social institutions such as "money, property, marriage, governments, elections, football games, cocktail parties and law courts" (Searle 1995, p. xi). Thus, for example, a certain kind of money is brought into existence as a new institutional fact when there is an agreement to ascribe the status of being money to some type of entity such as conch shells that previously existed only as a brute fact, and the form of this agreement is what Searle calls a constitutive rule, taking the form 'X counts as Y in context C' (Searle 1995, p. 28). The meaning of *status function* and *constitutive rule* is clear enough from this example. I have suggested some modifications to the concept of a constitutive rule elsewhere, but this does not alter the broad thrust of the argument (Elder-Vass 2012, chapter 4).<sup>1</sup> For the purposes of this paper, we must focus on the least clear of the three elements in Searle's argument, the concept of *collective intentionality*.

This concept appears in Searle's story because, on the one hand, institutional reality depends on the way that people think about it. Ten euro notes, for example, only work as money because (and when) people think that they are acceptable as money. The philosophical concept of intentionality captures this idea of mental properties that are *about* something. But on the other hand, an

---

<sup>1</sup> I have since discovered that Frank Hindriks had already developed a similar response to Searle's concept of constitutive rules (Hindriks 2009).

individual cannot make a ten pound note into money just by individually believing that it is so. There must be some sort of collective element for intentionality to do this job. Searle, like most collective intentionality theorists, is reaching for a concept of collective intentionality that can perform the role required to support institutional reality without implying the existence of group minds (Searle 2005, p. 21). Only individuals, he believes, can form intentional beliefs, and therefore if collective intentionality is real and significant, it must exist as the intentional beliefs of individuals but in some sort of collective form. Searle postulates that humans are biologically capable, as individuals, of forming “we-intentions” (Searle 1995, p. 24–26). Thus, for example, a violinist in an orchestra can believe that “we intend” to play a symphony, and if she then goes on to play her part in the performance, this reflects her holding at least two distinct intentions: the we-intention and also her individual intention to play her own part (Searle 1995, p. 23–25). Thus collective intentionality, or at least we-intentionality, does not exist for Searle as intentionality-of-the-collective but rather as intentionality-of-the-individual-on-behalf-of-the-collective (Elder-Vass 2012, p. 58).

We-intentionality in an isolated individual, however, is not sufficient to establish an institutional fact. If I, for example, believe that “we” regard bars of chocolate as money, this is not enough to make them so: they only become money if there is an *actual* we, an actual group of people that shares this belief, and in particular, an actual group of people that is therefore prepared to accept chocolate bars in payment, on the basis of the belief that they may subsequently use them in payment themselves. This does not just apply to money: all institutional facts depend upon the existence of a genuine collective, a group of people who each hold some variant of the collective intention, thus creating what Raimo Tuomela calls a “group-social fact” (Tuomela 2013, p. 220). This has been the theme of critiques of Searle by a number of scholars (e.g. Tuomela 2002; Meijers 2003). Searle has tended to neglect the necessity for an actual collective in his presentation of collective intentionality, but there are signs that he does accept the argument (Elder-Vass 2012, p. 65–67). He frequently talks about the need for collective agreement, acceptance, or recognition for the establishment of an institutional fact (Searle 1995, e.g. p. 39, p. 46). This, however, is a touch ambiguous, since he has also defined collective intentionality as a we-intention of an individual, so it is not entirely clear whether collective agreement, acceptance and recognition might also be taken to be purely individual properties. More recently he has argued that “the object or person performs its function only in virtue of collective acceptance *by the community* that the object or person has the requisite status” (Searle 2005, p. 7, emphasis added) thus seeming to recognise that we need both collective intentionality in individuals and the existence of a social group of individuals who share that collective intentionality before we can have

an institutional fact. Intriguingly, in his latest book he seems to make the equivalent point when discussing normative institutions without making it when discussing institutional reality (Searle 2010, p. 44–47, 156; Elder-Vass 2012, p. 66–67).

As far as institutional reality is concerned, Searle's ambivalence relates to the role played by language in his argument. His colleague Jennifer Hudin argued at the 2014 Cambridge/Berkeley workshop on critical issues in social ontology that we need a causal account of the world up to the mental level, but the mental level gives us language and once we have language then further features of the world require not causal but constitutive explanations. Searle developed the argument, saying that institutional reality is thus not an emergent feature but a constituted one, and as a result the appropriate focus of analysis becomes the logical structure of the linguistic forms that constitute it rather than a causal explanation. While this explains his failure to make a clear statement about the causal roles of social groups, it does not seem to leave him with a tenable account of institutional reality. As the chocolate-bar money example above makes clear, it is quite possible for an individual to hold a *we-intentional* belief in the logical form specified by Searle without this generating the corresponding institutional fact. Something more is required: a causal process in which an actual group of people holding the relevant belief develops, and a continuing set of causal interactions that sustains those beliefs. Institutional facts, I suggest, need causal and not just linguistic explanations, and in developing these we will need to recognise the causal power of social groups or entities. There is, in other words, an important absence in Searle's theory of institutional reality: a recognition of its dependence on the causal powers of groups. One purpose of this paper is to fill that absence by providing a critical realist account of this capacity. Although my interest is initially in how collective intentionality enables us to account for the causal power of some groups, those causal powers in turn plug a gap in Searle's account of institutional reality.

### 3 Causal Powers in Social Ontology

If it is accepted that we do need real groups of people to make Searle's theory of institutional facts work (though without introducing group minds), then his theory raises the question of what part groups play in social ontology and in particular the question of whether we must move beyond ontological individualism to explain institutional reality.<sup>2</sup> Although Searle has described himself as

---

<sup>2</sup> By *ontological individualism* I mean the belief that there are no collective social entities with causal significance and thus that all social events are caused by human individuals (and, perhaps, other non-social entities such as material things).

a methodological individualist, I understand that this was intended merely to signal his rejection of any concept of a group mind and does not necessarily entail a denial of the causal significance of all collective social entities (Searle 1997, p. 450; Barnes 2002, p. 251; Meijers 2003, p. 172–175, 178). As the previous section suggests, I share Lawson's view that his ontology implicitly depends upon the ascription of causal significance to larger social entities (Elder-Vass 2012, p. 65–67; Lawson 2012, e.g. p. 350). Searle himself offers us no justification for such an ascription, but a justification is readily to hand in the form of the second tradition of social ontology discussed in this paper: the causal powers tradition of critical realism.

Like collective intentionality theory, critical realism arose from within the philosophical tradition, but unlike collective intentionality it has been extensively adopted by social scientists (though more by those with a theoretical/ontological orientation than those with an empirical orientation). The core arguments come from the work of Roy Bhaskar; again these are well known so they will be covered very briefly here. Bhaskar argues from the inadequacy of positivist accounts of causality, which see it as nothing more than a constant conjunction of cause and effect in the realm of *empirical* experience. These accounts, he argues, are radically inconsistent with actual scientific practice, and such practice, at least in the experimental sciences, instead suggests a very different conception of causality. Experimental science does not simply observe events, waiting for regularities, but instead manipulates them by restricting interfering forces to isolate the action of particular mechanisms. This is only possible if *actual* events are co-determined by multiple interacting causal forces, some of which may be disabled in particular cases, and it only makes sense as a way of understanding the world if such forces continue to operate in the world outside the laboratory (Bhaskar 1975, p. 33). The need to restrict interfering forces, and the possibility of doing so, can only arise if causal forces are tendencies rather than producing exceptionless empirical regularities, and if causal forces are *real* potentials that may or may not be realised in particular actual cases depending on the other forces at work. Thus the very practice of science implies that we live in a world of multiply-determined events, which are produced by the contingent interaction of multiple causal tendencies. He calls these tendencies causal powers, and argues that every causal power is an emergent property of a thing, a property that itself depends on the characteristic structure of the thing possessing it (Bhaskar 1975, p. 50–52). I have filled out Bhaskar's argument by positing that the causal powers of a thing of any given type depend on the sorts of parts and the sorts of relations between them that are required to make a thing of that type, and that they are produced by processes of interaction between the parts (Elder-Vass 2010, chapter 3). We

may then employ Bhaskar's term *generative mechanisms* somewhat differently than he does himself – more, perhaps, in the vein of Bunge – to refer to these processes that produce the causal powers of a thing (Bunge 1997, p. 414; Elder-Vass 2010, p. 23). In this view, mechanisms themselves are not powerful, nor are they things with powers, but rather processes that occur within and amongst the parts of a thing and that generate the powers of that thing. Thus, to give a schematic example, the interaction between the parts of a hi-fi system generates a power of the system as a whole to reproduce patterns of sound. Here the hi-fi system is the bearer of a power, and the processes of interaction between its parts, including for example the reading of a signal from an optical disk by a photocell detecting reflected laser light, the conversion of that signal into an electrical pulse, the amplification of that current, the variation in magnetic power produced by the amplified current in the speaker coil, and the resulting movement of the coil and thus the speaker cone and thus air in its vicinity are the mechanism that generates this power.

It is not immediately clear how this model may be applied to the social world, and a number of critical realist thinkers have developed differing, though substantially overlapping, accounts of social structure and its causal significance (notably Archer 1988 [1996]; Sayer 1992, 2000; Bhaskar 1993; Archer 1995; Lawson 1997, 2003; Bhaskar 1998 [1979]). This paper assumes that the most productive way to theorise social structure is to see it in terms of the causal powers of social entities (Elder-Vass 2010) or communities (Lawson 2012). Both Elder-Vass and Lawson use these terms to refer to social groups or entities whose parts are people (and sometimes other material entities too), and which have emergent causal powers that depend upon the way in which these people and other parts are organised, or related to each other, within the group concerned. Because the concept of social structure is used in many different ways in the literature (see Elder-Vass 2010, p. 76–86), and the details of that debate are somewhat tangential to the argument of this paper, I avoid the term below and instead focus on the more clearly defined concept of a *social entity*.

Perhaps the most obvious type of social entity is organisations, which consist of people related in ways that are defined by their roles in the organisation. Each of these roles is a bundle of norms about how a person in a given position in the organisation should act, including both how they should act towards other members of the organisation and how they should act towards outsiders when they are acting on behalf of the organisation (Elder-Vass 2010, p. 153–157). The consequence is that such organisations are social entities with emergent causal powers, in the sense that they have powers that would not be possessed by the constituent members and other parts, even collectively, if they were not organised



into such an organisation.<sup>3</sup> We may well be able to explain such powers in terms of the interactions between the parts of the organisation concerned: by describing the generative mechanism that produces the causal power. But this does not alter the fact that the power would not exist in the absence of the organisation, even if the parts of the organisation did exist but in some other configuration. The consequence of these powers is that both members of the organisation and outsiders act differently (or *tend* to act differently) than they would in the absence of these relations (Elder-Vass 2010, p. 158–159). Thus, for example, a sales assistant employed by a retail company may sell something that belongs to the company, and a purchaser may hand over their money to the sales assistant in the (correct) belief that in doing so they are paying the organisation (Elder-Vass 2010, p. 173–174). Furthermore, organisations may have collective causal effects, for example an orchestra may be capable of producing harmonious music, which the members could not do without being organised into such a social entity.

Organisations, however, are only one possible kind of social entity with causal powers. Others may include, at one extreme of scale, temporary small groups such as queues that come together informally and yet interact in a way with causal significance, and at the other, market systems in which large numbers of people interact in ways that produce unintended systemic consequences such as economic growth, inflation, and cycles of boom and economic crisis. In each of these types of case, it is possible that there are social entities with causal powers that arise from the particular ways in which individual humans interact in entities of this type.

One implication of critical realist social ontology is that if this is so, there must be a mechanism at work in each case, a process of interaction between the members of the social entity that depends on the way in which they are organised in this kind of entity, and produces the powers of the entity concerned. We can strengthen the case for believing that a particular type of entity has a particular type of causal power if we can identify an empirically plausible mechanism that would indeed produce this power and that does depend on the type of relations that prevail in entities of this type. This paper considers the possibility that forms of collective intentionality might provide or contribute to such mechanisms. If they did, this would help to support arguments for the empirical value of *both*

---

<sup>3</sup> I call this a relational form of emergence. It is a weaker form of emergence than the form typically implied in the philosophy of mind and often associated, for example, with the work of C.D. Broad (McLaughlin 1992; Bedau 1997; Stephan 2002; Elder-Vass 2010, p. 28–33). However, I have argued that it is both strong enough to resist eliminative (or ontological) reduction and more plausible than the version discussed in the philosophy of mind (Elder-Vass 2014). It is also virtually indistinguishable from what Tony Lawson calls strong emergence (Lawson 2012, p. 351–353).



theories of collective intentionality *and* theories of the causal power of social entities.

Given that different social entities have different causal powers that depend on different mechanisms, collective intentionality might contribute to some of these mechanisms and not others, and it might contribute in different ways to different mechanisms. The challenge that we face in linking collective intentionality and causal power is thus not a single but a multiple challenge. This paper approaches that challenge by making a relatively detailed case for the argument that collective intentionality contributes to one such mechanism, and then opening out to consider a range of other potential cases much more schematically.

## 4 Norm Circles

The case to be explored in more detail is the claim that social norms – standards, guidelines or expectations about practices or behaviour – are a product of social entities called *norm circles* (Elder-Vass 2010, chapter 6; Elder-Vass 2012, chapters 2 and 3). The range of normative practices is vast, extending for example to the complexes of norms about sound patterns, meaning and grammar that constitute our languages, and those about what sorts of things it is acceptable to say and write. Norms are simpler social forms than institutional facts, indeed rules are a form of norm, hence given Searle's account of the role of constitutive rules it would appear that norms play an important part in underpinning institutional reality.

The significant claim of this theory is that norms are effective only because every norm is backed by a group of people: the norm circle for the norm.<sup>4</sup> The members of the norm circle for a given norm are those people who tend to

---

<sup>4</sup> Several participants at the ENSO III conference in Helsinki, October 2013 suggested that this account relates only to social norms, and that we may also have personal norms to which it does not apply. I can illustrate their point from my own experience: I do not eat mammals, whereas I do eat meat from fish and birds, but I am not aware of anyone else who follows this particular rule (I have encountered some chicken-eating vegetarians but they generally represent eating chicken as a failure to follow a social norm rather than an action in conformity with a norm or personal rule). I am not certain that it is appropriate to call such personal rules *norms* at all, but in any case they do not have corresponding norm circles. This, however, does not make my argument irrelevant to personal rules, or personal rules irrelevant to my argument: the personal rule in this example is derived partly from my experience of related social norms that do have norm circles, and innovations like this one feed into the social process of normative change. New norm circles grow up around some, but not others, of these innovations.

endorse and enforce it, for example by the use of praise, reward, criticism or punishment directed towards others on the basis of whether or not they conform with the norm. Norm circles are diversely intersectional: the membership of any one norm circle in contemporary society is likely to be different from that of any other, though these memberships overlap in diverse and complex ways. There are many further details of the argument but we can ignore them for the purposes of the current paper (for these details see Elder-Vass 2010, chapter 6; and Elder-Vass 2012, chapters 2 and 3).

The key point here is that norms are effective in influencing our behaviour, not because norms or rules as such have causal power, but because they are backed by social groups: norm circles. The mechanism is that, as a result of their experience of the endorsing and enforcing behaviour of members of the norm circle for a given norm, individuals understand that they face a social environment where people will reward or penalise them depending on whether or not they conform to the norm concerned. Thus, the influence of the norm circle is mediated through the beliefs or dispositions it tends to produce in individuals exposed to its influence. Indeed, the same is true of the members of the norm circle themselves, who in a sense feel they act on behalf of a wider force when they endorse and enforce the norm, again as a result of their past experience of the influence of the group. Hence both members of the norm circle and those exposed to its influence act differently than they would in the absence of such groups. As a result, norm circles tend to produce a tendency to comply with the norm.

There are some parallels between this discussion of norm circles and the earlier discussion of collective intentionality. On the one hand, we need to recognise the role of actual social groups if we are to explain normativity, just as we need to recognise the role of actual social groups if we are to make sense of the causal significance of collective intentionality. On the other, there is some approximation to collective intentionality at work within norm circles that contributes to the mechanism through which they are causally effective. But is it collective intentionality at all, and if so, what kind?

## 5 Collective Intentionality in Norm Circles

Debates in the literature about different kinds of collective intentionality have often taken a competitive form, with participants arguing that one form of collective intentionality is valid while another is not. There is an element of this, for example, in Searle's account of we-intentions, which are offered as a sounder basis for collective intentionality than the various formulations of common

knowledge, or “interlocking beliefs and expectations” (Ruben 1985, p. 107) as a condition of collective intentionality. This latter requirement implies that if two social actors both believe the same thing, e.g. that a ten euro note counts as money, this is not yet a collective intention, but rather only becomes one when it is also true that “x believes that y has these beliefs and expectations; y believes that x has these beliefs and expectations; x believes that y believes that x has them; y believes that x believes that y has them; and so on” (Ruben 1985, p. 109–110).<sup>5</sup> Ruben suggests that these various beliefs need only be “dispositional rather than occurrent” (Ruben 1985, p. 110), but this does nevertheless seem to require an implausibly complex cognitive structure to exist before we can count some set of beliefs as collective intentions. On these grounds, it might be argued that Searle’s we-intentions are a more plausible form of collective intentionality (Lawson 2012, p. 374–375).

However, different accounts of collective intentionality are not necessarily mutually exclusive: multiple versions of collective intentionality theory may be valid as descriptions of different types of empirical case. Thus, for example, Elisabeth Pacherie argues that some complex theories of shared intentionality may apply to adults but are too cognitively demanding to be attributed to young children, who nevertheless do exhibit a form of shared intentionality. She therefore draws on Bacharach’s work on team agency to develop a theory of “shared intention lite” that could plausibly be applied to children (and which seems quite close to Searle’s conception of we-intentions) (Pacherie 2013). This paper suggests that we do need theories of a variety of different forms of collective intentionality, not because of the differing capabilities of different individuals (though there may also be a case for that), but because collective intentionality may operate in different ways in different forms of social interaction.

Let us turn, then, to the question of how collective intentionality might contribute to the causal power of norm circles. I have suggested that a member of a norm circle will have a sense that in endorsing or enforcing the norm concerned, she is acting on behalf of some wider force than just herself, and that this sense of wider influence makes her more likely to back the norm than she would be without it. Indeed it is precisely this relational dependency that makes her action an implementation of the causal power of a norm circle and not just an individual act. This sense in itself arises from her previous experience of norm-endorsing and norm-enforcing behaviour by other members of the norm circle. Thus there is a kind of social dependency of her tendency to back the norm: she has a belief in a wider force, but that belief in turn arises from the existence of the wider force, the

---

5 Such accounts of joint intentions are still prominent (e.g. List and Pettit 2011, p. 33).

existence of the norm circle, and her experience of its empirical effects.<sup>6</sup> This does not mean that norm circle members necessarily have an accurate understanding of the nature of that force. If they were asked to explain their stance they might say, for example, that this wider backing came from God, or religion, or society, or the nation, or the law and it was their feeling that this particular force backed the norm that led them to feel justified in endorsing or enforcing it. But in each of these cases, the belief that the force concerned backs the norm is a product of their past interaction with individuals backing the norm. Or, they might believe more explicitly that it is other people, or some specific group of other people, that give wider backing to the norm and thus justify their own action in its support, but without having a very clear or accurate idea of *which* other people would back the norm.<sup>7</sup> There is a sense, then, in which their tendency to back the norm is a kind of *we*-intention, but it is a *we*-intention in which the individual's sense of the 'we' is potentially very vague and potentially inaccurate.

In any of these cases, however, a key factor is that the individual member of the norm circle has an expectation, not only that there is some wider force backing the norm, but also that others feel the same way and would therefore tend to back her up in endorsing or enforcing the norm. This is enough to give us something like David Ruben's "descending reason relations" (Ruben 1985, p. 113): the individual is influenced by the belief that others do or would back them in their action in support of the norm, and therefore that they can expect support for this action. These beliefs will tend to be true if there is a norm circle for the norm concerned and its members are at hand, and if this is not the case, these beliefs will soon be undermined by contrary experience.

If all this is sound, then each member of the norm circle (a) has a *we*-intentional disposition to act in support of the norm (to endorse and/or enforce it in her interactions with others), which means that, whether she has consciously verbalised this or not, she expects that her own disposition to act in support of the norm is involved with a more widely held, and thus at least implicitly collective, disposition to act in support of it; and each member (b) is right to expect this, even though she may have a very nebulous or inaccurate sense of who the other people are who might share that disposition. This is a minimal kind of collective intentionality: the individual need not have a clear sense of who the 'we' involved in their intention is, and may even misattribute some part of the sense

---

<sup>6</sup> It is also possible that there are cases where individuals mistakenly believe that a norm is backed by a wider group. They may still endorse and enforce it, but their tendency to do so is likely to decline if and when they realise that other people do not back them when they do so, or if and when they realise that other people do not endorse or enforce the norm themselves.

<sup>7</sup> See the discussion of imagined norm circles in Elder-Vass 2012, p. 127–130.

of collective identity entirely, e.g. to a god or gods; she need have no beliefs about other people's beliefs, except that she believes at least some of them are likely to support her in her endorsing and enforcing action; and she may even have quite a different internal understanding of the norm than others who endorse and enforce the same external behaviours.<sup>8</sup> Nevertheless, this also seems to be a strong enough form of collective intentionality to play a central part in the mechanism that produces a shared commitment to enforcing a norm among the members of a norm circle: and thus, strong enough to provide the emergent causal power of norm circles that underpins the effectiveness of all social norms. We may also trace a further link to Searle's version of institutional reality: institutional reality depends on our following the relevant constitutive rules, and according to the account of normativity above, we follow social rules because there are norm circles that cause us to tend to do so. Institutional reality depends on rules which depend on norm circles which depend on a weak form of collective intentionality (Lawson 2012, p. 375).

## 6 Collective Intentionality and Causal Power

Norm circles thus depend for their causal power on a kind of collective intentionality, but a kind that is much less complex, much less cognitively demanding, much less conscious, and much less homogeneous between the different participants than the kinds that are most commonly described in the philosophical literature. This does not entail, however, that we must discard all other collective intentionality theory in favour of this model. On the contrary, there may also be cases that correspond to other forms of collective intentionality theory. In order to give a flavour of this argument, this section will consider a range of other forms of social interaction and briefly comment on the extent and kinds of collective intentionality and causal powers that are implicated in them. (I use *forms of interaction* as a term of convenience here in order to reserve judgement, for a few lines, on whether any given form should be counted as a social entity with causal powers over and above those of the individual members.) The forms to be considered are:

- a) two people passing unacknowledged in the street;
- b) a queue;
- c) two people walking together;
- d) a business corporation;
- e) a market system.

---

<sup>8</sup> On this last point, see Elder-Vass 2012, p. 50–54.

a) First, let us consider the case of two people passing unacknowledged on the street. The case nicely illustrates the point that there is no form of social interaction without some basis in normative structures: the very practice of not acknowledging (while politely ignoring) someone unfamiliar that we pass in the street is common in many cultures (though by no means universal) and is firmly based in normative conditioning, particularly the kind of normative conditioning that tends to develop in urban cultures that must find ways to cope with the overstimulation produced by cities (Simmel 1971). Thus even this case depends on a kind of collective intentionality: the kind that occurs in the norm circle that underpins this norm. But the first significant question here is whether any *further* collective intentionality is involved in this practice, and there does not seem to be: we expect nothing from the other party beyond our privacy. And the second is whether this form of interaction generates new causal powers: as far as I can see it does not.

b) The case of a queue is more complex. A queue is a kind of temporary interaction group (Elder-Vass 2010, p. 146–149).<sup>9</sup> But here there seems to be rather more collective intentionality involved. When we join a queue we implicitly acknowledge that we have become part of a collective that shares certain intentional attitudes: we all individually wish to gain access to whatever resource the queue is for, we all accept that access to this resource must be serialised, and we implicitly undertake to join the queue, treat it as an agreed and accepted mechanism for achieving this serialisation, and follow the related set of queuing practices. In these respects, there is quite a high degree of homogeneity of intention (though it would be unusual for the members of a queue to have surfaced all of these intentions clearly to consciousness). The consequence of this set of practices is that we collectively form a social entity – the queue – with the causal power to serialise access to the resource concerned in a form that avoids most of the potential conflict over the process (Elder-Vass 2010, p. 146–149). In doing so, we create an institutional fact: a group of people only counts as a queue and is only able to perform the functions of a queue because we collectively recognise it as such. The ‘we’ is well-defined in this case: it is the intentions of the people who join it, and of those who interact with it (e.g. the people providing the resource being queued for), that make it a queue. All of this occurs, however, without any explicit agreement between the parties concerned.

c) The case of walking together has been theorised elegantly by Margaret Gilbert (1990). When two people agree to walk somewhere together, they create a collective intention. In agreeing to walk together, they implicitly agree that they will conform to certain standards that are normatively expected of people with

---

<sup>9</sup> Again, like all these examples, it depends on a social norm, or indeed a complex of related social norms, but from here onwards this will not be discussed: the focus of this section is on whether any *further* collective intentionality is involved beyond ordinary normativity.

such an agreement. This includes, for example, that they will keep pace with each other, and continue walking together until they have reached the agreed destination, unless they explicitly agree to terminate the agreement at some point. In such cases, mutual awareness of the commitments that have been made is high (though this depends on the unstated assumption that both share the same normative understandings of what is involved in walking together) and so this case approaches the model of collective intentionality that requires common knowledge of each other's intentions. We could perhaps argue that a group of people walking together are a social entity, but this is a social entity whose only causal consequence seems to be the effects it has on the members of the group: their sense of obligation to keep to their commitment.

d) One might imagine that a business corporation would be a paradigmatic example of full blown explicit collective intentionality in action. The employees of the business all know that it exists and that they work for it, they have explicitly agreed to do so and explicitly agreed the terms on which they will do so, the aims of the corporation are usually explicitly stated, and the responsibilities of individual employees are usually quite explicitly stated too. As a consequence of the employees interacting with each other and with outsiders in the ways specified in their roles, the corporation as a whole has emergent causal powers (Elder-Vass 2010, chapter 7). Thus we have a strong form of collective intentionality that makes a substantial contribution to the mechanisms underlying a clear set of causal powers. On the whole, this seems like a clear case in which common knowledge, explicit collective intentionality, and causal power all go hand in hand. Nevertheless, one might express a few doubts. Perhaps the most obvious is that it is unlikely that there is homogeneity of collective intentionality within most corporations of any size. Let us imagine that the corporation's official goals include profitability, growth, legality, and some level of integrity in dealing with employees and customers. In practice, we might expect senior executives to have bought into these goals (though even here there are exceptions) but it is not unusual for there to be many employees who do not share them all. An employee's intention in coming to work may be, for example, to earn a salary, enjoy working life as much as possible, and avoid getting fired, and for such employees there may be no internalised commitment to the apparent collective intentionality of the corporation at all. Such employees may be brought to contribute to the corporation's goals by a suitable system of disciplinary threats and rewards, but can we really say that in such cases collective intentionality is a significant mechanism underpinning the causal powers of the corporation? What seems to be causally significant is the acceptance by each position holder that they will individually perform their role in the corporation, but this need not entail the presence of any collective intentionality.



e) If we were to accept the implications of mainstream economics, one of the characteristics of market systems is a complete absence of collective intentionality. Individual optimising agents, on this view, pursue their own objectives with no concern for others' intentions, and engage in production and exchange in the market. The consequence, for this tradition, is that they tend to produce various systemic effects, such as price equilibration, efficient allocation of productive resources, and economic growth. From a more critical perspective we might modify this story in a number of ways. For example, economic actors are embedded in social networks and normative systems, they have non-individualistic concerns, and the consequences of market systems include unemployment, inflation, and cycles of boom and bust. But neither of these perspectives provides a collective intentionality of the market. Granted, in the heterodox perspective markets depend on normative practices, but even in this perspective there is no sense that economic actors create a market system by agreeing to treat something as a market system (though there are exceptions: e.g. formalised markets such as stock exchanges). Instead, a market system exists as an unintended aggregative consequence of large numbers of individual exchanges, and most market systems exist and have macro consequences independently of agreement about their institutional status and indeed often without any such agreement existing. This is not to deny that institutional facts can be built around market systems – perhaps initiatives to set up managed exchanges, or state regulation of aspects of the market, for example – but market systems as such are not institutional facts established by collective intentionality. Despite this, they do have causal powers: the power to produce at least some of the various systemic effects listed earlier in this paragraph.

These examples were intended to demonstrate the diversity of social forms, the variety of forms of collective intentionality at work within them, and the varying relationships between collective intentionality and causal power. Table 1 summarises the analysis.

**Table 1:** Forms of Interaction and Forms of Collective Intentionality.

Form of interaction	Form of collective intentionality	Causal power
Norm circle	Vague we-intention	Yes: social norms
a) Passing in street	None	No
b) Queue	Implicit but quite clear	Yes: serialising access
c) Walking together	Explicit, with common knowledge	Limited
d) Business corporation	Heterogeneous, varying from strong to virtually none	Yes: power of organisations
e) Market system	None	Yes: systemic economic effects

One might be forgiven for concluding that there is no relationship between collective intentionality and causal powers, but this is the wrong way to read this table. What it reflects is that some social causal powers depend on collective intentionality and others do not, and that the way in which a power depends on collective intentionality varies from type to type. The crucial issue for the purposes of this paper is whether the mechanism that underlies the power depends on collective intentionality, and what form of collective intentionality it depends on.

## 7 Implications

The collective intentionality tradition has paid relatively little attention to questions of causal power; and critical realists have paid relatively little attention to the question of collective intentionality. But this paper suggests that these two traditions could be brought together productively, and has attempted to illustrate the point with the case of norm circles. Norm circles are social entities with causal powers, and the mechanism that generates this causal power depends on a particular form of collective intentionality, though one in which the members' sense of the collective may be rather nebulous, and in which the content of the collective intention is limited to backing a particular social norm. Nevertheless, this collective intentionality ultimately depends for its causal significance on the existence of an actual collective composed of individuals who share the collective intention.

There are, however, many other forms of social entity with causal powers. It seems likely that some of these depend on stronger forms of collective intentionality than norm circles, whereas others may exert causal power without any dependence on collective intentionality at all. Understanding the relation between social causal power and collective intentionality more generally therefore requires us to recognise multiple forms of collective intentionality with differing roles in different types of social entity, as well as cases in which collective intentionality is at best peripheral.

Combining these two traditions therefore seems to be a productive endeavour, and an endeavour that begins to make the concept of collective intentionality relevant to sociological and causal-explanatory questions, and not just philosophical and analytical questions. Critical realism can supply the absent collective in Searle's account of social reality, and theories of collective intentionality help to explain at least some of the mechanisms behind the social causal powers that critical realists are interested in theorising.

## Bibliography

- Archer, Margaret S. (1988 [1996]): *Culture and Agency*. Cambridge: Cambridge University Press.
- Archer, Margaret S. (1995): *Realist Social Theory: The Morphogenetic Approach*. Cambridge: Cambridge University Press.
- Barnes, Barry (2002): "Searle on Social Reality: Process is Prior to Product". In: G. Grewendorf and G. Meggle (Eds.): *Speech Acts, Mind, and Social Reality: Discussions with John R. Searle*. Dordrecht: Kluwer Academic, p. 247–258.
- Bedau, Mark A. (1997): "Weak Emergence". In: *Philosophical Perspectives* 11, p. 375–399.
- Bhaskar, Roy (1975): *A Realist Theory of Science*. Leeds: Leeds Books.
- Bhaskar, Roy (1993): *Dialectic: The Pulse of Freedom*. London: Verso.
- Bhaskar, Roy (1998 [1979]): *The Possibility of Naturalism*. London: Routledge.
- Bunge, Mario (1997): "Mechanism and Explanation". In: *Philosophy of the Social Sciences* 27, p. 410–465.
- Elder-Vass, Dave (2010): *The Causal Power of Social Structures*. Cambridge: Cambridge University Press.
- Elder-Vass, Dave (2012): *The Reality of Social Construction*. Cambridge: Cambridge University Press.
- Elder-Vass, Dave (2014): "Social Entities and The Basis of Their Powers". In: J. Zahle and F. Colin (Eds.): *Rethinking the Individualism-Holism Debate*. Cham: Springer, p. 39–54.
- Gilbert, Margaret (1990): "Walking Together". In: *Midwest Studies in Philosophy* XV, p. 1–14.
- Hindriks, Frank (2009): "Constitutive Rules, Language, and Ontology". In: *Erkenntnis* 71, p. 253–275.
- Hund, John (1998): "Searle's *The Construction of Social Reality*". In: *Philosophy of the Social Sciences* 28, p. 122–131.
- Lawson, Tony (1997): *Economics and Reality*. London: Routledge.
- Lawson, Tony (2003): *Reorienting Economics*. London: Routledge.
- Lawson, Tony (2012): "Ontology and The Study of Social Reality". In: *Cambridge Journal of Economics* 36, p. 345–385.
- List, Christian and Philip Pettit (2011): *Group Agency*. Oxford: Oxford University Press.
- McLaughlin, Brian P. (1992): "The Rise and Fall of British Emergentism". In: A. Beckermann, H. Flohr and J. Kim (Eds.): *Emergence or Reduction?* Berlin: De Gruyter, p. 49–93.
- Meijers, Anthonie W. M. (2003): "Can Collective Intentionality Be Individualized?" In: *American Journal of Economics and Sociology* 62, p. 167–183.
- Osborne, Thomas (1997): "The Limits of Ontology". In: *History of the Human Sciences* 10, p. 97–102.
- Pacherie, Elisabeth (2013): "Intentional Joint Agency: Shared Intention Lite". In: *Synthese* 190, p. 1817–1839.
- Ruben, David-Hillel (1985): *The Metaphysics of the Social World*. London: Routledge and Kegan Paul.
- Sayer, Andrew (1992): *Method in Social Science*. London: Routledge.
- Sayer, Andrew (2000): *Realism and Social Science*. London: Sage.
- Searle, John R. (1995): *The Construction of Social Reality*. London: Allen Lane.
- Searle, John R. (1997): "Responses to Critics of *The Construction of Social Reality*". In: *Philosophy and Phenomenological Research* LVII, p. 449–458.
- Searle, John R. (2005): "What is an institution?" In: *Journal of Institutional Economics* 1, p. 1–22.

- Searle, John R. (2010): *Making the Social World*. Oxford: Oxford University Press.
- Simmel, Georg (1971): "The Metropolis and Mental Life". In: D. N. Levine (Ed.): *On Individuality and Social Forms*. Chicago: University of Chicago Press, p. 324–339.
- Smith, Christian (2010): *What is a Person?* Chicago: University of Chicago Press.
- Stephan, Achim (2002): "Emergentism, Irreducibility, and Downward Causation". In: *Grazer Philosophische Studien* 65, p. 77–94.
- Tuomela, Raimo (2002): "Searle, Collective Intentionality, and Social Institutions". In: G. Grewendorf and G. Meggle (Eds.): *Speech Acts, Mind, and Social Reality: Discussions with John R. Searle*. Dordrecht: Kluwer Academic, p. 293–307.
- Tuomela, Raimo (2013): *Social Ontology*. New York: Oxford University Press.
- Wettersten, John (1998): "The Analytical Study of Social Ontology: Breakthrough or Cul-de-Sac". In: *Philosophy of the Social Sciences* 28, p. 132–151.

---

**Article note:** Some of these arguments were presented at the ENSO III conference in Helsinki, October 2013, and the Critical Issues in Social Ontology workshop in Cambridge, May 2014. I would like to thank the participants for their stimulating and useful comments.