

Can Groups have Concepts? Semantics for Collective Intentionality

Cathal O'Madagain¹

Abstract: A substantial literature supports the attribution of intentional states such as beliefs and desires to groups. But a lacuna exists in this literature, which is the lack of much positive work on whether groups can have concepts. Since on many views, a subject cannot be attributed a belief without being attributed concepts, such a gap undermines the cogency of accounts of group intentionality. Here I draw on Lewis's conventional semantics to propose an account of group concepts, and argue that the same considerations put forward for the irreducibility of group intentional states to the intentional states of a group's members can be marshaled in support of an account of group concepts. I discuss a real-life example of a group concept, which we find in the case of the meaning of 'meter' as fixed by the International Bureau of Weights and Measurements, and finally I consider the upshot of these considerations for the question of social externalism about concepts.

1. Introduction

We have several reasons to think that the intentional states of groups form a distinct species in the classification of mental kinds. They appear prominently in our folk psychology, where we readily attribute beliefs and desires to governments, corporations, or simply groups of friends making plans for the weekend. They are not easily identified with the sum of the intentional states of the members of groups, since two groups with the same membership can be coherently attributed distinct beliefs (Gilbert 1989). And compelling group decision-procedures can result in a group endorsing a view that all of its members, individually, reject, suggesting a fundamental discontinuity of the group belief from the beliefs its members (Pettit 2001, 2003). Attributing intentional states to groups, however, may commit us to also attributing concepts to those groups, which on most theories of intentionality are the components of intentional states. But as yet we have very little by way of an account of group concepts. Here I propose such an account.

First I argue that something like Lewis's conventional semantics is necessary to assign meaning to the sentences groups use to express their judgments. I then argue that such a semantics has the same properties as group intentional states that make those states irreducible to the intentional states of group members, thus making the semantics similarly irreducible. To illustrate the point, I discuss a real-

¹ Thanks for helpful discussion to the editors of the current volume, to audiences at University College Dublin, Trinity College Dublin, University College Cork, Queen's University Belfast, Edinburgh University, to Philip Pettit, Deborah Tollefsen, Elisabeth Pacherie, Francois Recanati, Jonas Akerman, Sonia Sedivy and Bill Seager, and to Richard Davis of the *Bureau International des Poids et Mesures*. Research for this paper was supported by a Fernand Braudel Postdoctoral Fellowship, funded by the Fondation Maison des Sciences de l'Homme and the European Union Cofund.

life example of a term whose meaning seems to be fixed by a group rather than any individual – the term ‘meter’, which is set by the International Bureau of Weights and Measurements. Adopting a thin functional characterization of concepts, I argue that the meanings fixed by such a semantics should be understood as group concepts. And finally, I draw out some implications of these considerations for the question of social externalism about concepts.

2. From Group Intentional States to Group Concepts

Of the many accounts of group intentional states available (Gilbert 1989, Searle 1990, 1995, Pettit 1991, 2001, 2003, Tuomela 1995, Bratman 1999, Gilbert 2001, Tollefsen 2002, Bacharach 2006, Gold and Sugden 2007, Pacherie 2007, Tuomela 2007, List and Pettit 2011), ingredients that are common to most in one form or another are an individual intention condition, an interdependence condition, and a common knowledge condition. Assuming such a framework for simplicity, let’s suppose (roughly following List and Pettit 2011: 33-34)² a group *G* jointly takes an attitude to a proposition *p* if and only if:

- i) All members of *G* intend to do their part in bringing it about that the group believes/desires *p*
- ii) All members intend to do their part because they expect that all other members of *G* will do their part
- iii) i) and ii) are common knowledge in the group

Where *p* is a proposition – some ‘that’ sentence, for example ‘that Dr. No is the best candidate for the job’ – we can say the group believes that Dr. No is the best candidate for the job when the above conditions are met with ‘believes’ in condition i. In that case, each member intends to do her part in bringing it about that the group believes *p*, and each does so because the other members do too. Doing your part in bringing it about that the group believes *p* may involve endorsing that *p* whenever speaking on behalf of the group, acting as though *p* were true when acting on behalf of the group, and so on.

Conditions such as these do a great deal of explanatory work, including differentiating the kind of committees considered by Gilbert, for example. Suppose that of several standing committees in a philosophy department, the hiring committee and the Christmas party committee have the same members. Nevertheless, if the hiring committee decides on Dr. No for the job, something seems to have gone wrong if we infer that the Christmas party committee thereby hired Dr. No. The above conditions explain this. All members of the hiring committee may publically endorse the decision to hire Dr. No because they believe the other

² List and Pettit’s characterization of the conditions underlying joint belief/desire are, for the group’s members, that “they each intend that they together act so as to form and enact a single system of belief and desire [...]; they each intend to do their own part in a salient plan for ensuring group agency within that scope, believing that others will do their part too. And all of this is a matter of common awareness” (List and Pettit 2011: 34).

members of the hiring committee do, and the other members of the hiring committee may also be members of the Christmas party committee. But reasons are individuated intensionally – a reason that averts to the hiring committee is not the same as a reason that averts to the Christmas party committee, because there are possible scenarios where the membership of each committee is distinct. Once we include a *de dicto* reference to the relevant group in an account of its members' intentions, the intentional frameworks underlying the beliefs of the hiring committee and the Christmas-party committee can be distinguished.³

Such conditions also distinguish two groups that behave identically, but only one of which acts collectively, as discussed by Searle (1995)⁴. If a collection of individuals runs for shelter from the rain in a park, each member intends to execute the movements that get them, individually, to shelter. A dance troupe, on the other hand, may carry out a performance that requires its members to individually execute the same movements – and yet there is clearly an important difference between the dance troupe and the collection running from the rain. The interdependence condition captures this difference – only in the dance troupe do the members intend to execute the movements in question because they believe the other members of the troupe do too.

Although conditions such as these are considered by some to count as a reduction of group beliefs to the beliefs of their members (Bratman 1999), Pettit (2001, 2003) has shown that even with such a relatively simple model of group intentionality, groups can have intentional states that conflict with the views of even their entire membership, suggesting a significant 'autonomy' (List and Pettit 2011: 73-83) of group intentional states from those of their members. Suppose a group commits to a decision-making procedure according to which the group endorses any proposition entailed by a series of already endorsed propositions, and rejects any proposition ruled out by earlier endorsed propositions, in order to ensure consistency in the official position of the group⁵. A three-judge court, for example,

³ Although in Gilbert's (1989) account, the reference to the group appears in an initial commitment made by each member of the group to do their part in the shared plan of action – not quite the same as the interdependence condition included here, but closely related.

⁴ Although Searle's (1995) account does not appeal to these conditions to explain the difference, but rather to the notion of a primitive 'we intention' adopted by each member of the group. For difficulties with this approach see Baier, 1997; Stoutland, 1997; Velleman, 1997.

⁵ Such a procedure may be necessary for the group to be attributed the status of a rational agent (Pettit 2003, List and Pettit 2011), or may be adopted simply in order to ensure consistency in a group's decisions on a series of inter-related cases. Without adopting such a procedure, a majoritarian decision procedure can encounter the following difficulty. Suppose a group of three members, A, B and C, votes on three questions – p , if p then q , and q . If a majority is required to endorse a proposition, the group can vote as follows: a third endorses p (A and B), a third endorses if p then q (B and C), and a third rejects q (A and C). No member of the group displays inconsistency here, but the views attributed to the group are now inconsistent – a problem called the 'discursive dilemma' (Pettit 2001), which is a version of Condorcet's Paradox (Condorcet 1785), itself generalized as a problem for group decisions in Arrow's Theorem (Arrow 1951/1963). The problem can be avoided by the adoption of a 'sequential priority procedure' (Pettit 2001, List 2004, List and Pettit 2011) – where the group endorses only propositions that do not conflict with propositions already endorsed, and automatically endorses any proposition entailed by propositions already

may be faced with deciding whether a defendant is guilty of first-degree murder, which would require that the defendant caused the death of another, did so intentionally, and with malice aforethought. The following pattern of beliefs may hold among the judges A, B and C:

	Cause	Intent	Malice	Guilty?
A	Yes	Yes	No	No
B	Yes	No	Yes	No
C	No	Yes	Yes	No
Guilty?	Yes	Yes	Yes	?

Since none of the judges, individually, believe all conditions hold, no judge individually believes the defendant satisfies the conditions for being guilty of first-degree murder (the far right column). But if the court has decided to endorse any proposition that is entailed by already endorsed propositions, and it determines its view on the three 'premises' of the decision by majority vote, then the court will be compelled to rule that the defendant is guilty, since all premises receive a majority of positive votes given this pattern of beliefs among the judges. The view of a group can therefore depart radically from the views of its members.

We appear, as a result, to have a strong case for group intentional states, where those states are states of groups *per se*, and not of their members. If, however, we are to attribute beliefs and other intentional states to groups, are we required also to attribute concepts to groups? According to the representational theory of mind, the received view in the cognitive sciences (Pinker 2003) and much of philosophy (Carruthers 2000, Millikan 2000, Fodor 2003, Harman 1987, Margolis & Laurence 2007), intentional states are composed of parts which themselves represent objects and properties. We call these parts concepts. Distinct beliefs, for example, represent objects and properties as standing in different relations, due to the distinct concepts composing the beliefs, and their order in the belief. If Bond believes Dr. No is worse than Jaws, then Bond has a belief whose parts are the singular concepts DR. NO and JAWS, and the relational predicate concept IS WORSE THAN; the relationship between the parts in the belief determines whether the belief is true or false, depending on whether the fact represented by the parts appearing in this order obtains in the world or not. Some reject this picture and argue that concepts are abilities rather than representations (Dummett 1993, Kenny 2010) or even abstracta (Peacocke 1992). But on all of these accounts, it is necessary to possess or grasp the concepts that determine what you are thinking about, in order to be thinking about anything at all. As Fodor puts it, 'if a creature can think of something as green or triangular, then that creature has the concept GREEN OR TRIANGULAR. Further research is not required' (2008: 138).

In the case of group intentional states such as joint commitments, the question arises at the level of the semantics of the sentences the group uses to express their view. Invariably, when a group jointly commits to something, they do

endorsed. In the case considered, if the group votes first on *p*, then on if *p* then *q*, *q* will be endorsed automatically, and needn't be called to vote, thus ensuring against this kind of inconsistency.

so by jointly accepting the proposition expressed by some sentence in a language. When the court considered above is characterized as judging that the defendant is guilty of murder, this will inevitably involve their joint acceptance of the truth of the proposition expressed by some sentence, such as 'the defendant is guilty of murder'. Ordinarily, when an individual expresses their judgment by uttering a sentence, we attribute to the individual the concepts required to fix the semantics or meaning of the sentence. If Bond asserts 'Dr. No is worse than Jaws', we assume Bond grasps the concepts that give his utterance the meaning it has, or that fix the semantics of the sentence he utters. Similarly, when the court commits to the truth of the proposition expressed by the sentence 'the defendant is guilty of murder', we may seem required to attribute to the court the concepts that fix the semantics of that sentence.

To some, however, the attribution of concepts to groups seems *prima facie* problematic. Robert Rupert (2005), for example, gives the following consideration against the possibility of attributing concepts to groups. Our best naturalized theories of concepts, he holds, include indicator semantics (Dretske 1988), pure-informational semantics (Fodor 1987), teleo-semantics (Matthen 1984, Millikan 1984), causal-historical semantics (Rupert 1998, Prinz 2002) and teleo-isomorphic semantics (Cummins 1995). These accounts each involve various mechanisms that groups appear not to have. A central example he considers, on which I will focus, is perceptual states. In the theories under consideration these states either track, are causally correlated with, were originally caused by, were naturally selected to identify, or both were caused by and bear an isomorphic relation to, the objects or properties that a subject's concepts refer to. Rupert then points out that groups *per se* don't have perceptual states (and likewise for the other properties of individuals these theories appeal to). Therefore, he argues, groups *per se* cannot have concepts. And from there he concludes, reasonably, that there is something suspicious about the claim that groups *per se* can have intentional states.

I think that Rupert's argument is not valid, but before considering how to avoid the worry he raises, there is another important possibility to consider. Might group commitments come into being on the basis of a simple coincidence of the semantics assigned by individual group members to the sentence they consider? Perhaps, when the three-judge court expresses their judgment using the sentence 'the defendant is guilty of murder', there is nothing beyond the semantics assigned by each of the three judges in their own idiolects to determine the semantics of the sentence. On this view, the group *per se* would not need to be attributed concepts for the group *per se* to be attributed beliefs. This is problematic, however. As discussed, most current approaches to group beliefs involve an interdependence condition of some sort: when committing jointly to the truth of a proposition, each member endorses the proposition in part because the others do. These members must therefore not only entertain the same proposition, as argued, but they must also know that they are entertaining the same proposition. I can certainly endorse the truth of a proposition without knowing whether anyone else is entertaining the same proposition, but I cannot coherently endorse the truth of a proposition *because* others are endorsing its truth, without knowing that the others are endorsing the same proposition that I am. In other words, for a group to jointly

commit to some proposition expressed by a sentence 'p', the semantics for the sentence must not only be shared by the members of the group, it must be knowingly shared. A simple coincidence of the idiolects of a group's members will not, therefore, suffice to establish the semantics of the sentence jointly endorsed. As I explore next, the conditions required to ensure a semantics is knowingly shared in a group are sufficient to justify attributing that semantics to the group per se. And once we are in a position to attribute a semantics to a group per se, it is a small step from there to the attribution of concepts to the group.

3. Lewisian Semantics for Group Intentionality

What does it take for a semantics to be knowingly shared in a group? Lewis's (1969) account of linguistic meaning may have the answer. Lewis's account derives from his more general theory of convention, on which some regularity *R* in the behavior of a group is a convention if and only if:

1. Everyone conforms to *R*
2. Everyone expects everyone else to conform to *R*
3. Everyone prefers to conform to *R* on condition that the others do
4. It is common knowledge in the group that 1-3 hold (Lewis 1969: 59)⁶.

These conditions closely match the ingredients that have come to form the basis of most approaches to group intentionality, considered above. Conditions 2 and 3 can be collapsed into the single condition that corresponds to ii) above – that everyone conforms to *R* because they expect everyone else to (if they conform because they expect the others to, this entails that they expect the others to conform, and prefer to conform if others do). To integrate Lewis's model of convention with the views we have on group intentional states, then, I will represent Lewis's model in terms of the simpler three conditions.

Applied to language, Lewis's takes the notions of truthfulness and trust as the regularities in question. The idea is that a sentence 'p' means *p* in a language used by a group *G* if there is a convention of truthfulness and trust in *G* to use 'p' to mean *p*. Or, that a convention is in place to only say 'p' if *p* is the case, and conversely to believe that *p* is the case whenever someone says 'p'. Simplified to the three conditions considered so far, a language is conventional in a group *G* iff:

- i. Everyone in *G* is truthful and trusting in *L*
- ii. Everyone in *G* is truthful and trusting in *L* because they believe everyone else in *G* is truthful and trusting in *L*
- iii. i and ii are common knowledge in *G*

Might these conditions provide the semantics for the target sentences in group intentional states? There are two standard difficulties faced by Lewis's account as an

⁶ Lewis elaborates more complex conditions later on (1969: 78), and again in his (1975), to block subtle counterexamples. I will omit these complications here.

account of meaning in natural languages in general, which we might worry could get in the way. First, it is not clear that this convention has to be satisfied for ordinary uses of language in order for those uses to be effective. Davidson (1986) considers the case of Mrs. Malaprop, who uses language in bizarre and yet intelligible ways, lauding her guest, Captain Absolute, as ‘the very pineapple of politeness’ and hoping her daughter will make a ‘not altogether illegible candidate’ for marriage. Her audience understands her perfectly, and yet there cannot be a convention between Mrs. Malaprop and her audience to use these sentences with these meanings, since her audience has never considered using sentences in this way before. Davidson believes that the kind of *ad lib* interpretation required of the audience in this case is a pervasive component of natural language exchanges, which he thinks cannot, as a result, be based on convention. Developing a satisfactory account of linguistic meaning that respects the central role of such interpretation is perhaps the primary task of pragmatic accounts of linguistic meaning (Sperber and Wilson 1986, Carston 2002, Recanati 2004, 2010).

A second worry concerns the productivity of the language-faculty of individuals. We expect an individual’s language will determine a meaning for an infinite number of possible sentences – the indefinitely many compositional variants of the terms the semantics assigns meaning to. Speakers can have only a finite amount of evidence about the language-use of others, however, so it would seem that speakers can only establish truthfulness and trust in a finite number of sentences. As a result, it seems to some (Hawthorne 1990, 1992, Schiffer 1993, 2006), that the meaning of these indefinitely many sentences cannot be determined by conventions of truthfulness and trust to use the sentences in such ways. The meaning of each speaker’s utterances are determined, more likely, by their individual language faculty, which fixes a systematic and infinitely productive pairing of sentences with meanings, as Chomsky’s (1980) account holds.

However, neither of these worries arises, or arises as acutely, for the case of the sentences expressing group commitments. Consider Davidson’s worry. Mrs. Malaprop succeeds in communicating to her audience in unconventional ways, because her hearers are adept interpreters. But there is an asymmetric relationship between Mrs. Malaprop and her audience that is not admissible in a case of joint commitment. Although her audience knows that they don’t use the sentences Mrs. Malaprop produces in the way that she uses them, Mrs. Malaprop doesn’t know this – she thinks she is using these sentences in the typical way. Ordinarily, that doesn’t interfere with communication – it doesn’t ordinarily matter if you know your audience uses the sentences you produce in the same way that you do, or if they figure out what you mean on the spot. All that matters is that both speaker and audience assign the same meaning to the speaker’s utterances (Pagin 2008). But in the case of joint commitment, this does matter. As discussed above, because of the interdependence condition – that each party commits to the proposition under consideration *because* the others are also so committed – each party must know that the others understand the sentence expressing the joint commitment in the same way. A joint commitment could not, as a result, be established under conditions such as those in the scenario involving Mrs. Malaprop. While such exchanges may be a feature of natural language that Lewis’s model does not account for, then, they can’t

form the basis of a joint commitment – and so they don't rule out Lewis's account as a model for the semantics of joint commitment.

Hawthorne and Schiffer's worry, on the other hand, is that due to the infinitely many sentences in a language its semantics determines meaning for, the semantics cannot depend on its speakers' knowledge of other speakers' use of sentences, which will always be finite. This may show that an infinitely productive semantics for a language cannot be supported conventionally. But it does not show that no language can be supported conventionally. Rather, the most it shows is that any language supported by a semantics that does depend on convention will be finite in productivity – that is, that the semantics for such a language does not determine meaning for infinitely many sentences. Such a language might determine meaning for just one sentence – such as 'the defendant is guilty of murder'. Or, it might determine meaning for a finite number of combinations of the terms in such a sentence – such as 'the defendant is guilty of murder', and 'the defendant is not guilty of murder', where all the terms except 'guilty' and 'not guilty' in these two sentences are taken to have the same meaning in each case. But that is surely a tolerable limitation on a language that supports joint commitments. In order to commit jointly to the proposition expressed by the sentence 'the defendant is guilty of murder', the members of the court do not need to know, for infinitely many sentences, whether the semantics for those sentences is knowingly shared by other members of the court. They only need to know that the semantics for one sentence is knowingly shared – the sentence, 'the defendant is guilty of murder'.

Allowing that the classic difficulties for a conventional semantics do not, then, arise at the level of group commitments, how should we expect this semantics to fit into our account of group intentional states? A straightforward way is to introduce the semantics as a general precondition on the establishment of a group intentional state, which will inevitably involve the group deliberating on the truth of some sentence. Accordingly, the conditions required for a group *G* to take an attitude to a proposition *p* will be that:

- i) 'p' is a sentence in a language *L*, where:
 - a. Everyone in *G* is truthful and trusting in *L*
 - b. Everyone in *G* is truthful and trusting in *L* because they believe everyone else in *G* is truthful and trusting in *L*
- ii) All members of *G* intend to do their part in bringing about that the group believe/desire *p*
- iii) All members intend to do their part because they expect that all other members of *G* will do their part
- iv) i), ii) and iii) are common knowledge in *G*

In principle, *L* needn't provide the semantics for any more than the sentence 'p'. It may be the case that the members of the group assign wildly divergent semantics to sentences other than 'p', perhaps unknowingly (as in the case of Mrs. Malaprop), where those sentences are not used to establish the group intentional state. As long as *L* is sufficient to establish the meaning of 'p', the group can jointly take an attitude to *p*. Or the group may require *L* to be quite extensive, encompassing a wide range of

sentences that appear on the group’s agenda, and rules of composition for the terms in these sentences, allowing the group to extrapolate further commitments that the group is bound to, given previously endorsed commitments, to some finite degree. Either way, that L is a convention in the group, and ‘p’ a sentence in L, will be a condition on the group having a joint intentional attitude to *p*.

Supposing that such a semantics does underwrite the formation of group intentional states, do we have any reason to think of this semantics as being attributable to the group per se, as we do for thinking of the intentional states of groups? Importantly, the most striking reason for thinking of group intentional states as being in some way autonomous from the states of group members applies just as much in the case of the semantics of those states as it does for the states themselves. At least for isolated expressions, a group can explicitly deliberate on the meaning of an expression as it appears in group commitments, and the outcome of this procedure can exhibit the same features that lead us think of a group’s intentional states as irreducible to its members’. As an example, consider the procedure used to establish the meaning of the term ‘meter’ by the International Bureau of Weights and Measurements (BIPM). Due to increasing demand for higher precision in measurements by independent groups around the world who exchange, for example, mechanical parts for the auto or medical industry, or who combine the results of surveyors working in several parts of the world simultaneously, the BIPM has repeatedly refined the standard used as the reference of the term ‘meter’ since its inception in 1875. When the BIPM debate a new standard for the meaning of ‘meter’, they consider many factors, including whether a new candidate standard will be sufficiently continuous with the previous standard so as not to render useless records of earlier measurements made in terms of meters, whether the new standard will be an improvement in precision over the previous standard, and whether the new standard will be reproducible in different parts of the world, so that different groups will be able to invoke the common standard of measurement.

For example, the definition adopted in 1983, as the distance travelled by light in a vacuum in $1/299\,792\,458^{\text{th}}$ of a second, was adopted in part because widely available technology at the time made it possible to check exactly what that distance amounts to in many places around the world. This makes it useful for mass production in different countries of precise engineering parts, and so on. As it happens, the resolutions on the meaning of ‘meter’ have always been unanimous, but the BIPM allows that what is required to pass a motion is not in fact unanimity, but an absolute majority (a majority of all eligible voters) (BIPM 2013, §2.1.3.5). According to these guidelines, then, the BIPM could have found themselves with the following pattern of votes, on the questions whether the proposed standard was sufficiently continuous with the previous standard, sufficiently precise, and sufficiently reproducible, where A-C each represent a third of the votes cast:

Voters	Continuous	Precise	Reproducible	‘Meter’?
A	Yes	No	Yes	No
B	No	Yes	Yes	No
C	Yes	Yes	No	No

'Meter'?	Yes	Yes	Yes	?
----------	-----	-----	-----	---

As considered before, if the group adopts to endorse any proposition entailed by propositions already adopted by majority vote, then given this pattern of beliefs on the premises, the group will be compelled to adopt the new standard as the meaning of 'meter' even though no individual member of the voting committee was convinced of its adequacy. The meaning of 'meter' as it appears in subsequent statements made by the BIPM, and any group or individual deferential to the BIPM in its use of that term, will as a result display just the same kind of autonomy from the minds of the group members who contributed to establishing the convention as the verdicts of the three-court judge considered above display from the opinions of the individual judges. The semantics assigned to statements made involving this term is, then, not the product of an individual language faculty, but the product of a group effort. Since groups are free to adopt such procedures to determine the semantics of the sentences expressing their commitments, this shows that such a semantics displays just the same sort of autonomy from the idiolects of a group's members as group intentional states display from the intentional states of a group's members.

The central question of this paper, however, was whether groups can have concepts. So far I have argued that the meanings of the sentences expressing group commitments requires a conventional semantics, and that a conventional semantics can confer an autonomy on these meanings equal to that which we have found in the intentional states attributable to groups. But are these meanings, even if fixed by facts about groups *per se*, concepts? Rupert worries that our best theories of concepts make it implausible that a group might have concepts, since they appeal to features of individuals that groups don't have, such as perceptual states. But even if our current theories of concepts all involve an appeal to features of individuals that groups don't in fact have, would that mean that groups couldn't have concepts? Groups would be disbarred from having concepts, presumably, if our notion of a concept *essentially* involved a property of some sort or another that groups can't have. Our concept CONCEPT, however, hardly appeals to perceptual states essentially. First, it's worth noting that the role played by perceptual states varies considerably across the different theories Rupert considers. Perceptual states may play an essential role in the individuation of concepts on a causal-historical model of concepts such as that of Rupert (1998) or Prinz (2002). But in Fodor's informational semantics, although perceptual states play a role in what he calls a concept 'locking on' to its referent, famously the perceptual state plays no essential role in the individuation of the concept, which exists with its reference already determined prior to this event.

Taking a slightly wider sweep of theories of concepts still, we find the inferential-role theory of concepts (Harman 1987, Brandom 1994), which holds that concepts are individuated according to the inferences they license from thoughts they appear in, and makes no essential reference to perceptual states. And we also find a Platonic theory of concepts (Peacocke 1992), which takes concepts to be abstracta that can exist without ever being entertained in a mind, not to mention

interacting with perceptual states. The fact that these theories are not *a priori* false shows that perceptual states need play no role whatsoever in a theory of concepts. Rather, our notion of a concept would seem to have a much thinner, functional definition. The only essential role that concepts are required to play, common to all of these theories, is that of fixing the semantics of the semantic variables in intentional states. Accordingly, we cannot rule groups out from having concepts because they don't share with individuals the states that instantiate concept possession in those individuals. All that is required is that the candidate intentional states of groups can be assigned a semantics, and that this semantics is fixed by facts about the group per se. As I have argued, a semantics is available for group intentional states which does indeed seem to be fixed by facts about groups per se. The semantic variables that result, such as the conventionally fixed meaning of 'meter', should therefore be considered group concepts.

4. Group Concepts and Social Externalism

Lastly I want to consider what bearing these considerations have on our treatment of the question of social externalism. The phenomenon that originally motivated the idea of social externalism has come to be called 'deference' – the idea that in our use of some words and concepts, we appear to accept the authority of others on their content. The classic example is that of a patient who announces that he has arthritis in his thigh (Burge 1979). An individualist's theory of meaning holds that an individual's concepts pick out either what he would describe as falling under his concept (Russell 1912), or what causes his concept to token in his thoughts (Dretske 1981, Fodor 1987). And yet it would seem that this patient will give a false description of arthritis, as a condition that can occur in the thigh, and his concept is caused to token by something other than arthritis (the pain in his thigh). So it would seem that according to individualism, the patient couldn't be thinking of arthritis. Since ordinary practices of belief attribution will attribute to this patient a belief about arthritis, albeit a false one, Burge concluded from these considerations that individualism about mental content must be wrong. Burge suggested, alternatively, that the patient's grasp of the concept ARTHRITIS depends on his immersion in a linguistic community that includes experts who fully grasp the concept, and from whose competence the patient's ability to entertain a thought about arthritis somehow derives. On this view, the patient's grasp of the concept depends on his external social environment.

An initial response to this puzzle on the part of the individualist was to try to show that perhaps the patient really is not thinking of arthritis, but instead has his own concept THARTHRTIS, which picks out the condition in his thigh (Fodor 1987, Crane 1991). But this deprives us, problematically, of the intuitive attribution to the patient of a false belief about arthritis – rather than the patient's announcement that he has arthritis in his thigh showing that he has a false belief about arthritis, the view implies the deeply counter-intuitive result that his utterance expresses a true belief about tharthrtis. A more attractive treatment of the puzzle, and perhaps more widespread at this point, is to hold that our concepts can piggyback on the linguistic use of others in our community. When Burge's patient uses the term 'arthritis',

making claims about aches in his thigh or elsewhere, what he really means is 'whatever the experts around here call 'arthritis'' (cf. Chalmers 2002: 617). This approach allows that Burge's patient can have false beliefs about arthritis – he can be trivially wrong about what the experts in his community use the term 'arthritis' to refer to. This analysis is attractive to individualists because it keeps the individual in a place of privilege in our theory of concepts. Fodor, who eventually came around to the idea of deference, emphasizes this point: 'If my ELM concept is deferential, that's because the botanist's isn't' (Fodor, 1998: 154). The individualist's response to the challenge of social externalism, then, is to argue that the correct analysis of these cases involves a privileged role for individuals, in whose minds concepts must be ultimately 'grounded' (Recanati 2000).

The current discussion, however, suggests that this meta-linguistic analysis, if true, needn't entail a privileged status for the individual. If the term 'meter', for example, is used deferentially, then it seems that it is not true that there is a single mind at the end of the relevant chain of deference. Suppose two physicists have a dispute about the length of a meter. One says to the other:

- I. I think a meter is longer than $1/300\,000\,000^{\text{th}}$ of the distance travelled by light in a second

And the other replies:

- II. I disagree, I think a meter is shorter than $1/300\,000\,000^{\text{th}}$ of the distance travelled by light in a second

Presumably they each expect that the meaning of 'meter' as it appears in both utterances is settled independently of either speaker, or there wouldn't be much point in disputing the question (if the terms expressed idiosyncratic tharthritus-like concepts, they would both be right). And presumably, as physicists, they will both defer on the answer to this question to the BIPM. Let's suppose the term 'meter' in each utterance is therefore equivalent to 'whatever the BIPM has accepted as the meaning of 'meter''. To resolve their dispute, then, the physicists turn to the BIPM. What they will find there, however, is not one individual mind that settles the question. What the BIPM mean by 'meter' is determined by vote, the outcome of which needn't be identical with, and hence is distinct from, the opinions of any of its members. At least in this case, then, a chain of deference is not grounded in the opinion of an individual, but of a group.

5. Conclusion

I have argued here that a theory of group intentionality needs to be supplemented with an appropriate semantics. That the semantics for the job is that of Lewis (1969), which although problematic in the case of natural language in general is ideally suited to account for the content of group intentional states. That the reasons we have for considering group intentional states to be autonomous from the states

of group members apply just as much in the case of the semantic variables in those states, and that given a thin functional definition of concepts, those resulting variables should be understood as concepts. Finally, I argued that these considerations combined suggest that thinking about collective intentionality may provide new support for the view that in some contexts, the concepts that appear in the intentional states of groups or individuals may themselves also be fixed by facts about groups, not individuals.

References:

- Arrow, K. (1951/1963). *Social Choice and Individual Values*. New York: Wiley.
- Bacharach, M. (2006). *Beyond Individual Choice: Teams and Frames in Game Theory* (edited by N. Gold and R. Sugden). Princeton: Princeton University Press.
- Baier, A. (1997). Doing things with Others: The Mental Commons. In Alanen, L., S. Heinämaa, and T. Wallgren, eds., *Commonality and Particularity in Ethics*, New York: St. Martin's Press, Inc.
- BIPM. (2013). Compendium of main rules and practices applicable to the BIPM. Online: <http://www.bipm.org/utils/common/documents/official/compendium.pdf>
- Brandom, R. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*, Cambridge, MA: Harvard University Press.
- Burge, T. (1979). Individualism and the Mental. In French, Uehling, and Wettstein (eds.) *Midwest Studies in Philosophy IV*, University of Minnesota Press, pp. 73-121.
- Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Chalmers, D. (2002). The Components of Content. In *Philosophy of Mind: Classical and Contemporary Readings*, ed. D. Chalmers. Oxford University Press.
- Chomsky, N. (1980). *Rules and Representations*. New York: Columbia University Press.
- Condorcet, M. d. (1785). *Essay sur l'Application de l'Analyse a la Probabilité des Décisions Rendues a la Pluralité des Voix*. Paris.
- Crane, T. (1991). All the Difference in the World. Reprinted in S. Goldberg and A. Pessin (1996) *The Twin Earth Chronicles*. New York: M. E. Sharpe. 284-304.
- Cummins, R. (1996). *Representations, Targets, and Attitudes*. London: MIT Press.
- Davidson, D. (1986). A Nice Derangement of Epitaphs. In D. Davidson (2005), *Truth, Language and History*. Oxford: Clarendon Press. 89-108.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge Mass: MIT Press.
- Dummett, M. (1993). *Seas of Language*, Oxford: Oxford University Press.
- Fodor, J. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. London: MIT Press.
- Fodor, J. (1994). *The Elm and the Expert: Mentalese and Its Semantics*. London: MIT Press.
- Fodor, J. (1998). *Concepts*. Oxford: Oxford University Press.
- Fodor, J. (2008). *LOT2: the Language of Thought Revisited*. Oxford: Oxford University Press.
- Gilbert, M. (1989). *On Social Facts*. Princeton: Princeton University Press.

- Gilbert, M. (1997). What is it for us to intend? In *Contemporary action theory*, vol. 2, ed. G. Holmstrom-Hintikka and R. Tuomela, Dordrecht: Springer; 65–85.
- Gilbert, M. (2001). Collective Preferences, Obligations, and Rational Choice. *Economics and Philosophy* 17: 109–120.
- Gilbert, M. (2004). Collective Epistemology. *Episteme* 1: 0-17.
- Gold, N. and R. Sugden (2007). Collective Intentions and Team Agency. *Journal of Philosophy* 104: 109–37.
- Harman, G. (1987). (Nonsolipsistic) Conceptual Role Semantics, in E. LePore (ed.), *New Directions in Semantics*, London: Academic Press.
- Hawthorne, J. (1990). A Note on ‘Languages and Language. *Australasian Journal of Philosophy* 68: 116–118.
- Hawthorne, J. (1993). Meaning and Evidence: A Reply to Lewis. *Australasian Journal of Philosophy* 71: 206–211.
- Kenny, A. (2010). Concepts, Brains, and Behaviour. *Grazer Philosophische Studien*, 81 (1): 105–113.
- Lewis, D. (1969). *Convention*. Cambridge: Harvard University Press.
- Lewis, D. (1975/1983). Languages and Language. Rpt. in *Philosophical Papers*, vol. 1. Oxford: Oxford University Press.
- List, C. (2004). A Model of Path-Dependence in Decisions over Multiple Propositions. *American Political Science Review*, 98(3): 495–513.
- List, C. and Pettit, P. (2011). *Group Agency: the Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- Margolis, E. & Laurence, S. (2007). The Ontology of Concepts — Abstract Objects or Mental Representations? *Noûs*, 41 (4): 561–93.
- Millikan, R. (1984) *Language, Thought, and Other Biological Categories: New Foundations for Realism*. London: MIT Press.
- Pacherie, E. (2007). Is collective intentionality really primitive?. In M. Beaney, C. Penco & M. Vignolo (Eds.), *Mental processes: representing and inferring*, Cambridge: Cambridge Scholars Press: 153-175.
- Pagin, P. (2008). What is Communicative Success? *Canadian Journal of Philosophy*. 38(1): 85-115.
- Peacocke, C. (1992). *A Study of Concepts*, Cambridge, MA: MIT Press.
- Pettit, P. (2001). Deliberative Democracy and the Discursive Dilemma. *Philosophical Issues* (supp. to *Nous*), 11: 268–99.
- Pettit, P. (2003). Groups with Minds of Their Own. In F. Schmitt (ed.), *Socializing Metaphysics*, Oxford: Rowman and Littlefield.
- Pinker, S. (1994). *The Language Instinct: The New Science of Language and Mind*. London: Penguin.
- Prinz, J. (2002). *Furnishing the Mind: Concepts and Their Perceptual Basis*. London: MIT Press.
- Recanati, F. (2000). Deferential Concepts: A Reply to Woodfield. *Mind & Language*, Vol. 15 No. 4 September 2000, pp. 452–464.
- Recanati, F. (2004). *Literal Meaning*. Cambridge: Cambridge University Press.
- Recanati, F. (2010). *Truth-Conditional Pragmatics*. Cambridge: Cambridge University Press.
- Rupert, R. (2005). Minding One’s Cognitive Systems: When does a Group of Minds

Constitute a Single Cognitive Unit? *Episteme*. 1 (3):177-188.

Schiffer, S. (2006). Two Perspectives on Knowledge of Language. *Philosophical Issues* 16: 275–287.

Searle, J. (1983). *Intentionality*. London: Oxford University Press.

Searle, J. (1990). "Collective Intentions and Actions." In *Intentions in Communication*, P.Cohen, J. Morgan, and M.E. Pollack, eds. Cambridge, MA: Bradford Books, MIT press.

Searle, J. (1995). *The Construction of Social Reality*. New-York: The Free Press.

Sperber, Dan and Deirdre Wilson, 1986, *Relevance: Communication and Cognition*, Oxford: Blackwell.

Stoutland, F. (1997). 'Why Are Philosophers of Action so Anti-Social?' in Alanen, Heinamaa, and Wallgreen eds. *Commonality and Particularity in Ethics*. New York, N.Y.: St. Martin's Press.

Tollefsen, D. (2002). Organizations as true believers. *Journal of Social Philosophy* 23 (3): 395–410.

Recanati, F. (2000). Deferential Concepts: A Reply to Woodfield. *Mind & Language*, Vol. 15 No. 4 September 2000, pp. 452–464.

Recanati, F. (2004). *Literal meaning*, Cambridge: Cambridge University Press.

Recanati, F. (2010). *Truth conditional Pragmatics*, Oxford: Oxford University Press.

Russell, B. (1912). *The Problems of Philosophy*. Oxford: Oxford University Press.

Velleman, J. D. (1997). How to Share an Intention, *Philosophy and Phenomenological Research* 57: 29-50.