

## **Reply to Decker on ‘Coordinated Behavior, Emergence, and the Explanatory Salience of Collective Representations’**

Whit Schonbein  
College of Charleston  
schonbeinw@cofc.edu

[This is an elaborated version of my spoken replies given to Mark Decker’s comments at the Pacific APA meeting, April 10<sup>th</sup>, 2009. Please see the original paper (the title of which is contained in the title of this note) for more information.]

First, I want to thank Dr. Decker for his helpful comments. I think the points he brings up are for the most part correct. However, I think that several miss their targets. Here I’ll first explain why before turning to what I take to be the most substantive objection.

To briefly recapitulate, on my reading, Huebner’s initial objection is that if Rupert is right about the superfluity of group representations, then individual representations are also explanatorily superfluous, because they are both compositionally derived from constituent representations. I proposed that Rupert could avoid Huebner’s initial objection by reformulating his criteria for explanatory salience. Specifically, I proposed that the relevant condition actually had two parts, and summarized this as principle (R):

(R) An appeal to a representation is explanatorily superfluous when (1) there is an explanation of how the content of the representation is a function of its constituent contents (the compositionality constraint), and (2) the causal efficacy of the posited representation is explained by appeals to constituent representations of the same type as that representation (the causal constraint).<sup>1</sup>

I went on to (initially) apply (R) to two cases – that of a club deciding to order pizza, and that of a man, Ted, bursting into tears at the sight of his recently deceased mother. The point of the first case is to show that, if we take the club to have a group representation (e.g., a decision to order pizza), then both clauses of (R) are satisfied, and the representation is therefore explanatorily superfluous. The point of the second case is to show that, if we take Ted to have a MOTHER representation, the first clause of (R) is satisfied, but the second is not, and therefore the MOTHER representation is not superfluous. In this way Huebner’s initial argument is refuted – being derived is not sufficient for superfluity, on this interpretation of Rupert.

Decker discusses each of these cases; here I’ll reply to his comments on the pizza case, saving the case of Ted and his mother until later. Decker objects that Huebner would probably simply concede that the club did not have a non-superfluous group representation. I think this is right – there is no reason to believe that a proponent of collective representations would not simply allow that there is no explanatorily salient

---

<sup>1</sup> This version of (R) is slightly amended from the original version, as there were two typos in the original that rendered it rather difficult to comprehend. I thank Dr. Decker for not picking the low-lying fruit.

group belief or desire in that case. (Indeed, given Huebner's discussion of court decisions, I suspect that he would also agree that there is no such representation in the scenario that Rupert uses to motivate (R) – a court decision.) However, my purpose in introducing the pizza case was simply to illustrate how (R) functions as an alternative interpretation of Rupert's criterion for explanatory superfluity. Insofar as Hueber's reductio relies on a different (incorrect) interpretation, it is refuted, regardless of whether he takes the club case to be a genuine case of collective representation. Furthermore, I *do* go on to argue that (R) applies in the same way to a case where Huebner does claim there is a non-superfluous use of collective representation, that of ship navigation. So I think I can go ahead and concede Decker's point about the pizza ordering scenario without sacrificing my argument.

A similar reply holds for Decker's objection to my discussion of the ship navigation case. In the original paper, I try take Huebner's emphasis on coordinated behavior in ship navigation seriously by interpreting the scenario as an instance of a form of emergence as described by Andy Clark (1997), the hope being that such a form of emergence will be sufficient to ground claims of explanatory salience. To motivate this interpretation, I follow Clark's discussion, which focuses first on nest-building by termites, and second on the very same ship navigation research (i.e., Hutchins' work) invoked by Huebner. Emergence in the case of termites is motivated by the fact that, despite their capacity to build a nest, there is no individual-level representation of the nest – there is no full plan of the nest represented in any individual termite – so the capacity to build a nest 'emerges' from lower-level capacities. The strategy, then, is to apply this same lesson to the case of ship navigation, with the goal being that the putative collective representation of the ship's location as instantiated by the ship's crew is analogous to the nest-building capacity of the termites – it is emergent from individual-level representations and the activities of crew members.<sup>2</sup>

Decker objects to this strategy on the grounds that the termite and ship navigation cases cannot legitimately be compared. In particular, he doubts that there really is no member of the ship's crew that lacks an individual-level representation of the plan for acquiring a chart position of the ship's location, or at least a representation of that goal. This would appear to spell trouble for my appeal to emergence.

There's a lot that can be said in defense of the comparison between the termite and ship navigation cases. For example, it really *is* the case that no individual has complete knowledge of the plan for acquiring the ship's location, as Clark notes in describing Hutchins' research on ship navigation (Clark, 1997, p. 77; Hutchins, 1995) – the process is simply too complicated for any one individual to fully represent.<sup>3</sup> Furthermore, while

<sup>2</sup> Here is another way at viewing this strategy: I argued that principle (R) prima facie applies to the case of ship navigation, and hence an appeal to a group-level representation of the ship's location is explanatorily superfluous. But the applicability of (R) depends on the group representation being compositionally derived from individual level representations. So, perhaps this form of emergence will grant a non-compositionally derived form of higher-level representation (where being compositionally derived is interpreted more broadly than is allowed by Fodor). I don't think it does – instead, it shows that there is no higher level representation at all – for reasons paralleling the argument in the original paper.

<sup>3</sup> And furthermore, there is no need for any one individual to know the entire process. Instead, the overall process is distributed across many individuals, in the same way the process of nest building is distributed across many termites – or so the comparison goes.

it is the case that individuals (especially those officers who need a chart position to make decisions) *do* represent the *goal* of acquiring a chart position, aside from initiating the process, this goal does not play a role in actually generating the chart position. One could conceivably replace each navigation crew member with a dedicated bit of computational technology lacking any representation of the overall goal of its activities. Instead, like a robot in an automobile assembly line, a signal arrives causing it to execute its program, and it proceeds to acquire a bit of information, sending the output along some channel (the technology need not even have a representation of where it is sending the information, aside from an address of some sort). One robot might drop a plumb line off the bow and report the depth; another might drop a line off the stern, again reporting the depth; another might take a measurement off a visually salient geographical feature and report the distance; and so forth. The goal, as represented by the captain, may play a role in initiating this process – e.g., by causing the captain to press a button labeled ‘acquire location’ – but that representation is not part of the location acquiring system, and plays no role in guiding the steps in the process of acquiring a location fix.

However, even if the termite and ship navigation cases are not comparable, this would not amount to a defense of Huebner’s position. My goal in bringing in Clark’s notion of emergence was to provide a positive interpretation of Huebner’s argument – that is, I was trying to defend on Huebner’s position by bringing in additional theoretical apparatus to elaborate on his claims about ship navigation. If Clark is not allowed to compare the termite and ship navigation cases as he does in his (1997), then Huebner is not allowed to help himself to Clark’s notion of emergence in the case of ship navigation (assuming, of course, that he wanted to). So, again, I don’t see any danger in conceding Decker’s point (although, as mentioned above, I do think Clark is correct in comparing the two cases, and that Huebner can help himself to Clark’s notion of emergence if he wants to).

Turning, then, to the case of Ted and his mother. My claim in that case was that, unlike the case of the club ordering pizza, Ted’s MOTHER concept was still explanatorily salient despite being compositionally derived (in some non-Fodorian sense<sup>4</sup>) from the contents of its constituent representations. This is because, I claimed, the scenario failed to satisfy the causal condition of principle (R):

“We don’t explain Ted’s behavior by saying he had some set of visual representations plus an associated affective representation. Rather, we say he broke down in tears because he recognized the deceased person in front of him *as his mother*. It is the fact that the [representation] has *that* content that explains Ted’s behavior.”

In reply, Dr. Decker objects that we can certainly explain Ted’s behavior by appealing to the representations and processes that comprise his MOTHER concept. And if such an explanation is perfectly legitimate, then the addition of the causal clause to (R)

---

<sup>4</sup> Non-Fodorian because such a semantics would appear to violate the stricture that the content of molecular representations be a ‘straightforward’ function of the contents of constituent representations along with their contents.

does not defeat Huebner's original objection that Rupert's argument implies that individual representations are explanatorily superfluous.

I must agree with Dr. Decker in his initial point: Yes, it could be a legitimate explanation of Ted's behavior to appeal to the representations and processes underlying his MOTHER concept. However, I think that the causal condition of (R) still renders the MOTHER concept explanatorily salient in a way that the decision of the club to order pizza is not – for reasons that were unfortunately unclear in the original paper.

To motivate this point, first let me be explicit about what I take to be the 'constituent-representation' explanation we are granting may be explanatorily appropriate in Ted's case. Such an explanation would appeal only to 'sub-personal' representations such as states of the limbic system, states of the temporal lobe (e.g., areas that represent the presence of eyes, noses, lips, hairlines, and their relationships), and other such states that comprise Ted's MOTHER concept. This, I believe, is in keeping with Huebner's claim that the MOTHER representation is composed of states of various parts of the brain (2008, p. 102).

Second, consider the following observations: (1) In Rupert's original argument – which involves explaining the causal role of a court decision – the explanation of that causal role involves appeals to the beliefs and desires of the judges as well as those of individuals affected by that decision. (2) In Dr. Decker's objection to the Ted scenario, he consistently refers to Ted's beliefs, desires, and other intentional states. These include Ted's loving his mother, and that if his mother were gone, he would miss her. Furthermore, Decker invokes the MOTHER concept, noting that the sub-personal representations either constitute or cause that concept.

These observations suggest that, while Ted's behavior *could* be explained at the level of 'sub-personal' representations, there is also an explanatory context that requires explanations at the level of beliefs and desires. In particular, the relevant context for (R) is belief-desire explanation (an observation I take to be implicitly supported by Rupert and Decker's choice of language, as described in the previous paragraph). One widely accepted constraint on such explanations is that the causes they identify are also *reasons*. That is, the explanations identify the causes of behavior, but they also make sense of that behavior.<sup>5</sup> So the causal condition in (R) is asserting that the causal role of a compositionally derived group representation must function as a reason for the group's behavior in a way that the constituent representations do not.

So, consider the 'sub-individual' explanation of Ted's crying, which explains the behavior by informing us that light struck his retina, activated certain areas in visual cortex, which activated representations of facial features and their relationships in the temporal lobe, which in turn were associated with areas of the limbic system that, when activated, caused Ted to cry.<sup>6</sup> To paraphrase Dretske (1988), such an explanation would explain why the *crying* happened, but it would not explain *why* the crying happened. To

<sup>5</sup> Here I have in mind philosophers such as Quine, Davidson, Dennett, and Dretske, but the list could be vastly expanded. Of course the sense in which propositional attitudes function as reasons differs greatly between the various views, but the point is merely that attitudes do function as reasons.

<sup>6</sup> Here I am assuming there is no 'mother' cell, because the point is to avoid invoking the MOTHER concept in explaining Ted's behavior. If there is a candidate MOTHER representation in this explanation, then it is not clear that we have an explanation at the level of constituents as required by Huebner.

explain the latter, we need to be made aware of why it makes sense (/is rational?) for Ted to be crying in that situation, and to make that point, we need to invoke the MOTHER concept.<sup>7</sup>

If this is correct, then principle (R) implies that a necessary condition on the non-superfluous explanatory use of a group representation is that the group be the sort of thing for which propositional attitudes can function as reasons for behavior. One popular proposal is that, for this to work, an entity must be subject to standards of rationality – a proposal developed by Philip Pettit (2007) with respect to group behavior. However, it should also be noted that it seems unlikely that the ship navigation case satisfies Pettit's constraints. Which brings me back to the conclusion of the original paper – even if a group has an emergent state (in Clark's sense), that in itself will not be enough to guarantee the salience of collective representations – additional (as yet unspecified, but perhaps along the lines proposed by Pettit) conditions must be met.

---

<sup>7</sup> We of course also need to invoke other beliefs and propositional attitudes, as Dr. Decker points out.