

Collective Understanding

Delarivière, Sven

Published in:
AVANT

DOI:
[10.26913/avant.2020.02.01](https://doi.org/10.26913/avant.2020.02.01)

Publication date:
2020

License:
CC BY-NC-ND

Document Version:
Final published version

[Link to publication](#)

Citation for published version (APA):

Delarivière, S. (2020). Collective Understanding: A conceptual defense for when groups should be regarded as epistemic agents with understanding. *AVANT*, 11(2), 1-20. <https://doi.org/10.26913/avant.2020.02.01>

Copyright

No part of this publication may be reproduced or transmitted in any form, without the prior written permission of the author(s) or other rights holders to whom publication rights have been transferred, unless permitted by a license attached to the publication (a Creative Commons license or other), or unless exceptions to copyright law apply.

Take down policy

If you believe that this document infringes your copyright or other rights, please contact openaccess@vub.be, with details of the nature of the infringement. We will investigate the claim and if justified, we will take the appropriate steps.



Collective Understanding—A conceptual defense for when groups should be regarded as epistemic agents with understanding

Sven Delarivière

Centre for Logic and Philosophy of Science
Vrije Universiteit Brussel
Sven.Delariviere@vub.be

Received 29 August 2018; accepted 19 July 2019; published 6 January 2020

Abstract

Could *groups* ever be an understanding subject (an epistemic agent ascribed with understanding) or should we keep our focus exclusively on the individuals that make up the group? The way this paper will shape an answer to this question is by starting from a case we are most willing to accept as group understanding, then mark out the crucial differences with an unconvincing case, and, ultimately, explain why these differences matter. In order to concoct the cases, however, we need to elucidate what it means to be attributed with understanding and what makes up an epistemic agent. I shall argue that it is abilities, above all, that guide our attributions of understanding. While it is true that understanding must go beyond single acts, this is not going beyond as in going behind them (to private occurrences which are impossible to discern and don't themselves contribute anything of value), but beyond as in considering what people could and would do: their discernable and valuable abilities. To conceptualize whom the abilities belong to, I will specify what it means to be an epistemic agent. I shall argue that it is being a successful target of the epistemic stance. The epistemic stance (heavily inspired by Dennett's intentional stance) forms an instrumental abstraction by attributing epistemic properties (i.e., beliefs, epistemic aims, problem-solving tactics) to an entity to explain or predict its behavior. Macro-systematicity (a higher-level pattern which a theory can exploit) is what makes the epistemic stance's abstraction an explanatory powerful one (regardless of how that macro-systematicity is realized) and emergence (the lack of a straightforward mapping-relation between the micro and the macro) is what makes its power unique to a particular level of explanation. To end, I suggest two kinds of mapping-relations and walk through the shift from reducibility to emergence.

Keywords: cognition; collective understanding; group agents; epistemic stance; emergence

Introduction

The concept of ‘understanding’ has recently started to gain ground in epistemology and the philosophy of science. What has not received an equal amount of attention in this literature is how to conceptualise who understands.¹ I would like to contribute to the explication of ‘understanding subject’. Traditionally, philosophers have taken for granted that a relevant entity ascribable with understanding must always be an individual human. Could other entities be worthy candidates for understanding-attributions, though? The question of this paper is whether *groups* could ever be ascribed with understanding. As such, the aim of this paper is to answer the following Research Question, which will be abbreviated to (RQ):

(RQ): (i) What would make groups an understanding subject (an epistemic agent ascribed with understanding) and (ii) what should make us change the subject from a group’s member(s) to the group collectively?

Before we shape an answer to (RQ), it is important that we first further specify what the question means (in section 1). First, I will characterize what it means to understand. Here, we will keep our focus on a paradigmatic example of an understanding subject, namely an individual human expert, and consider what it is about this example that makes us ascribe or deny understanding. I will argue that it is abilities, above all, that guide our attributions. That groups can display abilities is uncontroversial enough, but whether it is appropriate to attribute those abilities to an epistemic group agent is another question. To understand this question, I will go over the concept of epistemic agency and consider what it is about individual human experts that would have us ascribe them with such agency. Then we can see whether the same applies to groups or whether there are any worthwhile differences that would have us abstain from considering groups in a similar light.

The way this paper will shape the answer to (RQ) is by starting from what we are most willing to accept as group understanding and work our way backwards. I concoct (in section 2) the two most extremes: (a) a conceptually poor case of collective understanding and (b) a conceptually powerful (even if practically implausible) case of collective understanding. We can then mark out what the crucial differences are between the two which make us change our mind and subject, as well as explain why these differences matter. I will argue (in section 3) that it is macro-systematicity (i.e., a higher level pattern which we can detect and which a theory can exploit) and, I will elaborate (in section 4) that is those which are emergent (i.e., where the higher level ‘pattern’ has no straightforward mapping-relation to lower level ‘patterns’) which would have us attribute groups, and not its members, with understanding.

1. Characterizing Understanding & Epistemic Agency

(RQ) is a question with no easy answer, but we can make a lot of headway by starting with a clear and useful conceptualization. For that purpose, we will be concocting the most conceptually powerful case of *collective* understanding and consider its relevant differences from the least powerful case. This does entail that the crucial difference between the powerful and the

¹ One exception is Toon (2015).

poor case of collective understanding should lie in the kind of collective subject and not in the quality of understanding. Ideally, we keep the latter constant, but to do so we need to be able to express what it is that is being held constant and what is being varied. We need to elucidate what makes us (1.1) attribute understanding and (1.2) consider the target entity as an epistemic agent.

1.1. Understanding

It is fair to say that understanding is a trait that we sometimes uncontroversially attribute to certain epistemic agents. But is there a systematic way to reveal what understanding-attributions are or should entail? Answering this question involves specifying what we find so philosophically or epistemically valuable about understanding and thus necessary for its attribution, regardless of who it is attributed to or what it is about. What is it in those situations of uncontroversial attribution or dismissal that guides us in ascribing or denying understanding? With inspiration drawn from Ryle (2000), I shall argue that understanding-attributions always boil down to the subject in question possessing a particular set of relevant abilities.

Certainly, the most salient feature of understanding is the feeling or *sense* of understanding. While this sense may be interesting from a psychological perspective, it has rightly been criticized (e.g., Trout, 2002) as an unsatisfying characterization for understanding as an *epistemological* concept. It is not difficult to think of examples of someone understanding something without having any accompanying feeling and even easier to think of examples where someone did experience that sense but was simply overconfident.

The literature makes frequent references to understanding, involving ‘grasping’ (e.g., Kvanvig, 2003; Khalifa, 2013; Grimm, 2014), or ‘seeing’ (Zagzebski, 2001), but what is being grasped or seen and what that entails is something that varies depending on who uses the term, meaning it still requires spelling out before they characterize anything. Philosophical conceptions of understanding are invariably connected to the presence of mental states (be it as explanatory concept or trait) and abilities (be it as trait or symptom). So, the two main lines of spelling out the trait of understanding are as (a) appropriate mental states, potentially with abilities as symptoms (e.g., Zagzebski, 2001; Van Camp, 2013; Wilkenfeld, 2013) or (b) appropriate abilities² (e.g., de Regt & Dieks, 2005; Ylikoski, 2009; Hills, 2015; Delarivière & Van Kerkhove, 2017). Siding with the latter, I will argue why considering abilities as a mere symptom of mental states would put the cart in front of the horse and create needless problems. The upshot is that it is not the mental states themselves that are empirically accessible or epistemically valuable to us, so we both detect and judge mental states by the abilities, and not vice versa. Allow me to elaborate.

If we characterize understanding through mental states, then the successes of an understanding subject will need to come from the appropriate mental occurrence, namely that subject considering something appropriate in their mind’s eye or other mind-organ (i.e., having the appropriate mental state) and the abilities that the subject may display are merely the public

² Potentially with certain mental states as explanatory constructs, as we shall explore later with the Epistemic Stance.

effects of the appropriate private mental occurrences. This view fits with some of our everyday language about understanding and justifications for understanding (e.g., seeing how things fit together, pointing to gaps or flaws in a mental image). There are, however, a couple of problems: First, if understanding is characterized through the appropriate private occurrences, then the task of gauging another person's understanding is not just difficult, it is impossible. We cannot look into anyone's mind, private as it is. We have no idea whether Euclid actually understood geometry, for only Euclid himself knew what went on in the privacy of his own mind. So, either our attributions of understanding are vain exercises in metaphysics or attributions of understanding are not really about these secret occurrences. The second problem is that there's no clear way to decide which occurrences are appropriate by referring only to those occurrences themselves. What is the correct mental state to understand the irrationality of the square root of 2? This is not only hard to characterize, but the avenues in which we look to justify our choices are invariably motivated by the external affordances granted, not by the intrinsic value of the occurrence itself. Calling the occurrence in one's mind's eye 'a proof' would be vacuous if what is seen cannot be put to any public use (or worse yet, the wrong use). Conversely, if we found out that someone exhibiting extreme competence had mental states different from what we have hitherto characterized as appropriate, we would broaden the scope of appropriate mental states, not deny them with understanding. One way out of this problem would be to say that the right occurrences are those of which the public effects are certain appropriate abilities, a route taken by Wilkenfeld (2013),³ for instance. But then our understanding-attributions are decided by the abilities and not the occurrences which supposedly lie behind them. We may as well put the horse in front of the carriage, where it can do its work unencumbered.

One reason why it may still seem appealing to refer to private occurrences beyond a person's acts is because it is possible for the same act to be performed with and without understanding, e.g., by sheer luck, by rote imitation, or by blind rule-following. This seems to imply that the difference-maker for understanding lies not in the act, but in something beyond it. But we may readily concede this without thereby having to withdraw into a secret world. The use of the plural in "abilities" is not incidental. Understanding involves more than just having a single ability appropriate to the object of understanding. I can easily memorize a correct response to a certain question (or even a few of them) without understanding what it is that I am saying. The problem with such a single-track disposition or ability is not that such an ability should be discredited, but that the scope of abilities is too narrow. Answering questions according to a set script may be something one can get away with in tests that happen to only involve the memorized questions (giving a misleadingly good impression), but this will quickly fail once tested for other appropriate abilities (e.g., correct courses after a setback, explain in different words, apply it in a practical circumstance, make an analogy, answer what-if-things-had-been different questions, etc). It is not the narrow *success* of memorization that is problematic, but the wide *failing* it entails that makes for a poor understanding. If we draw open the scope of understanding to include all the relevant abilities, then it becomes increasingly difficult to

³ "[O]ne's understanding x consists in large part of representing x in the right sort of way (...) I will contend that 'the right sort of way' is best cashed out as a mental representation the possession of which enables certain abilities" (Wilkenfeld, 2013, p. 1002)

motivate why a subject displaying the wide range of relevant abilities would *not* merit an understanding-attribution. When we attribute people with understanding, we are not making untestable inferences to any secret phenomena which are forever out of our reach and judgement, but we are gauging the appropriate multi-track dispositions of an understanding subject. So, it is true that our understanding-attributions go beyond single acts, but this is not going beyond as in going behind them (to occurrences which are impossible to discern and don't themselves contribute anything of value), but beyond as in considering what people could and would do, namely their abilities (which we can discern and value). We may, instrumentally, speak of mental states, but we do so on the basis of abilities, and not the other way around. Hence, it is the abilities that actually characterize understanding. As such, we can characterize understanding in the following way:

'S understands X' corresponds to 'S possesses sufficient abilities appropriate to X.'

The concept of abilities could be further cashed out as multi-track dispositions, but to keep this paper specific in focus and down in length, we will not be cashing out abilities beyond this. Which abilities are 'appropriate' to understanding has a variety of candidates, depending on who you ask and what kind of understanding they are meant to capture, including: recognizing qualitatively characteristic consequences (de Reg & Dieks, 2005), making counterfactual inferences in the contexts of manipulation, prediction and explanation (Ylikoski, 2009), answering what-if-things-had-been-different questions (Grimm, 2014), evaluate explanations (Khalifa, 2013), relating knowledge to other knowledge (Van Camp, 2013), etc. I believe the scope of understanding to be quite wide (i.e., composed of, but not exhausted by any single set of candidates discussed in the literature) and contextually dependent on the many aims and values of the epistemic practice to which the understanding-attributions relate. Having characterized understanding, even if roughly, we can further specify our initial question of this paper by replacing 'understanding' in (RQ) with 'possessing the appropriate abilities' in (RQ'):

(RQ'): (i) What would make groups an understanding subject (an epistemic agent that possesses the appropriate abilities), and (ii) what should make us change the subject from a group's member(s) to the group collectively?

It becomes relevant to note that we nonetheless do allude to mental properties in our explanations and descriptions of someone's abilities. We talk about someone having an appropriate belief, intending to achieve a certain worthwhile epistemic aim and having a sound tactic to do so. Furthermore, these ascriptions can be tied together into one explanatory entity to which the understanding can be attributed. From the entity's multi-track dispositions, we do not just note its abilities, but we discern that there is an epistemic agent to whom the abilities belong. This entails that a failure of collective understanding will either be due to a lack of abilities or a failure to attribute the abilities to an epistemic group agent.

1.2. Epistemic Agency

What is it that warrants anyone (or anything) being an epistemic agent? My contention is that it is nothing more or less than being a successful target of what I call the epistemic stance. This approach to epistemic agency is heavily lifted from Dennett's (1990, 2009) interpretationist conception of agents through the intentional stance. The intentional stance is a popular approach in the philosophy of mind. It has also been utilized successfully in social ontology by Tollefsen (2002, 2015) and List & Pettit (2006), although my account will depart from both of them at different stages of the paper. Here is a good summary of the intentional stance:

Anything that is usefully and voluminously predictable from the intentional stance is, by definition, an intentional system. The intentional stance is the strategy of interpreting the behavior of an entity (person, animal, artifact, whatever) by treating it as if it were a rational agent who governed its 'choice' of 'action' by a 'consideration' of its 'beliefs' and 'desires.' (Dennett, 2009, p. 339)

The intentional stance is an innate capacity (as opposed to an academic theory) to interpret an entity as being governed by beliefs, intentions, and rationality. The sole justification for considering an entity as an agent is the efficacy of the stance in predicting or explaining its behavior that way (regardless of how it is realized physically), so there's no difference between a 'real' and an 'as if' agent, and no dividing line between the two (Dennett, 2009). It is a normative stance in that this interpretation depends on what the agent ought to do (rationally), and requires a holistic approach in that the success of the stance lies not in pairing up the components of the stance with particular behaviors, but in how well the agent-package predicts or explains the entity's behaviors overall (Dennett, 1990).

In the same line, an epistemic agent is nothing more or less than the successful target of the intentional stance, except with a focus on the epistemically relevant properties. To distinguish the two, I shall refer to *the epistemic stance* instead, but the difference with the intentional stance is not one of kind, but one of focus. The epistemic stance is the strategy of interpreting the behavior of an entity by treating it as if it were governed by beliefs (which are always epistemic), epistemic aims (i.e., the kind of results that an epistemic practice values), and epistemic tactics (i.e., any serious systematic attempt to get closer to an epistemic result), as well as any other intentions that play a supporting role in the epistemic agency. Like the intentional stance, it is normative and the normative standards of an epistemic practice are brought to bear in the interpretation. What counts and what does not as epistemic will not always be neatly distinguishable. Keeping a room at a particular temperature is not an epistemic aim. Applying for funding or drinking tea to help with concentration is only on the very fringes of epistemic aims or tactics, whereas breaking down the problem into parts or writing down a mnemonic device to help navigate a search-space is properly epistemic. I shall assume that the reader and I share similar enough intuitions about what counts as proper epistemic and what does not at all, and leave the grey areas for what they are. Given the intuitive agreement of many thought-experiments in epistemology, it seems that if the epistemic stance is not innate, then it is certainly deeply rooted.

The stance is a holistic and instrumental abstraction. It is instrumental in that the sole justification for interpreting an entity as an epistemic agent is the explanatory and predictive success of that interpretation. If seeing an entity *as* governed by epistemic attributes (beliefs, epistemic aims, and tactics) has explanatory or predictive power, then, by definition, that entity is an epistemic agent. We can now further amend our (RQ') by specifying that epistemic agenthood is warranted by being a successful target of the epistemic stance in (RQ'')

(RQ''): (i) What would make groups, displaying the appropriate abilities, ever be a successful target of the epistemic stance, and (ii) what should make us change the subject (the relevant entity to be attributed with understanding) from a group's member(s) to the group collectively?

Both the concept of understanding and epistemic agency start from the multi-track dispositions of the entity in question. Furthermore, one tends to invite the other. It would be unlikely that an entity displaying a set of abilities does not allow for an explanatory story of beliefs, epistemic aims, and tactics accounting for its actions. And, conversely, it would be unlikely to find an entity with no appropriate abilities whatsoever, but where an interpretation of beliefs, epistemic aims, and tactics does, nonetheless, have explanatory value. That groups can display abilities is uncontroversial enough, but the question remains whether it is appropriate to attribute those abilities to an epistemic *group* agent, and this question is now in an appropriate shape for us to answer it in the remainder of the paper.

2. The Two Extremes

We will now concoct the least and most conceptually powerful case of *collective* understanding to use them in shaping our answer to (RQ).

2.1. The Composite Class

The two cases need to meet some requirements to ensure that we are not misled by the wrong kinds of differences between them. First, they both need to involve a collective of people in a way that they can be labelled as a "group" (in the broadest sense of the term), for the only potential candidates of collective understanding are, not surprisingly, collectives. Second, both need to be tested for their abilities (understanding) in the same way, so the difference lies with them, not their tests. For this, let's assume they both get the same written exam. Let's furthermore assume that this exam is set up with incredible care, such that many of the relevant abilities (appropriate to the object of understanding) are tested for and that it would be difficult to cheat or take shortcuts (e.g., by memorizing a few answers to standard questions), meaning that the results of the exam are a good indicator of understanding. Last, the results of the candidates in both cases should be equally successful. What makes the least conceptually powerful group the least conceptually powerful should not be that it does not display the relevant abilities in any way, for then the attribution of collective understanding would fail simply because there is no understanding to be attributed, and not because it was not collective. So, our two candidate groups need to display, in some form or another, the *same* relevant abilities for understanding,

but have it come from entirely *different* kinds of groups. With all this in mind, I submit Composite Class as the *least* conceptually powerful case:

Composite Class: Arbitrarily chosen people (including experts and non-experts) are put in a class and given one exam to fill in. Each class-member in turn adds their own answers to the exam.

What makes the Composite Class a ‘group’ is that the people chosen are strung together under the label (and roof) of a ‘class,’ and made to contribute to the same exam. So, what makes them a group is nothing more than being a literal composite of the individual members and their exam consists of nothing more than fragmented results, which are a literal composite of the individual answers—some of which are correct and some of which are not. While the Composite Class does satisfy the condition that abilities need to be present (some of the answers being those of experts), we have strong reasons not to talk about understanding belonging to the group, except as a shorthand to talk about (some of) its individuals. To employ the epistemic stance towards the group as a whole would be very difficult and most unrewarding. The class does not act like any epistemic agent we know, nor will its exam results even resemble that of any epistemic agent. Seeing the class as a literal sum of *several* epistemic agents will explain the exam and/or predict future exams but trying to see the exam as the output of a single epistemic agent would inevitably fail because the exam is simply too fragmented to benefit from an extra epistemic stance fixated on the group as a whole.

2.2. *The Expert Planet*

For the *most* conceptually powerful case of a group understanding, I would like to draw heavy inspiration from a paradigmatic individual epistemic agent: a human expert on the topic at hand. With this, I do not wish to make the claim that a human expert (especially not any particular one) should be the ultimate and singular standard for what characterizes an understander, but if I am secure in the claim that a typical human expert is a good example of one, then I can also say that any functionally equivalent entity would be an equally powerful example of an understander. Gilbert’s (2013) does justice to this idea when she says a ‘group belief’ entails that the members ‘are jointly committed to believing that p as a body’ (p. 137), which requires them ‘to emulate, as far as possible, a body that believes p’ (p. 140). With this (and other reasons that will be explored soon) in mind, I submit the Expert Planet as the most conceptually powerful case of a group understander or a group epistemic agent:

Expert Planet: Every citizen on the planet takes on the roles of a single neuron such that the planet is isomorphic to the brain (and relevant body parts) of an expert. The planet is presented with a vast version of the exam and fills it out with its enormous planet-hands.

Some readers may recognize this case as similar to the China Brain (Block, 1978), except that the object of understanding is not Chinese, but any area of expertise covered by the exam. While there are a couple of other differences between the Expert Planet and the China Brain, none of them are intended to distinguish the two cases in kind. Instead they are small amendments intended to help guide our intuitions more appropriately: I have changed the nation into

a planet so that we are less focused on our intuitions about the individuals and politics of real nations. I have also added ‘and the relevant body parts.’ This is not just to account for the planet interacting with an oversized exam-paper, or to make it easier to picture the planet as a full-fledged entity, but also to take into account any possible cognitive roles the body might play in individual human cognition (see the debate on embodied cognition). To satisfy proponents of embedded and extended cognition, I am also willing to further amend the thought-experiment to include an oversized environment as it will not change anything about my argument.

The Expert Planet was set up to be functionally equivalent to a paradigmatic individual expert even in its most fine-grained detail. This means that every kind of behavior or disposition (from which we discern the abilities, beliefs, epistemic intentions and tactics) will be equally present in the Expert Planet. The beliefs, epistemic intentions and tactics which we use to predict or explain its behavior will be of equal use in predicting or explaining that of the Expert Planet. And because the epistemic stance makes no dictates on implementation, it does not matter to the stance that the entity is comprised of individuals rather than neurons. In short, everything that made that expert fit for employing an epistemic stance can also be found in the Expert Planet and every reason we have to treat the individual expert as an epistemic agent would apply for the Expert Planet. Nonetheless, it is a highly implausible scenario to expect to occur. So implausible that one may conceivably question why any ramifications of such a thought-experiment should be relevant to us. The reason why is because if there is any kind of group that would constitute a powerful case of group-understanding at all, it would be a group that, as a whole, operates *exactly*, even in its most fine-grained detail, like an individual expert. One of Rupert’s (2011) objections to group cognitive states is that it is very unlikely that there is any fine-grained similarity between the functional profile of individuals and any existing groups. But the Expert Planet is set up so that it definitely does. This makes its ludicrousness a strength. If you are not swayed to attribute collective understanding even in such an extremely unlikely situation as the Expert Planet, then I am not sure why any other collective effort would. Nonetheless, it does bear mentioning that one of the arguments surrounding the China Brain thought-experiment is exactly whether it is a convincing case of understanding (be it collective or not). Block (1978) has famously argued against functionalism on the basis that it sounds ludicrous to attribute the Chinese nation with mental properties. However, I will say that (i) the scenario’s practical implausibility may contaminate what we think about its conceptual appropriateness (i.e., planets or nations tend not to behave even approximately like an individual expert—which is not the same as their being unsuited to be deemed one if they would) and (ii) attributing the Chinese Nation or Expert Planet with these properties is consistent with our best explanation of individuals (functionalism about other minds and understanders) and it is not clear why Block’s unease with the idea should count for anything more destructive to functionalism than being a surprising entailment. At the very least, it shifts the burden of proof to the opposition. If one wants to insist on the Expert Planet being a case without any collective epistemic agency or understanding, one would need to clarify what the relevant difference is and why that difference is the one that matters.

We now have two cases, both of which display the successes of understanding, where one would benefit from an epistemic stance (towards the group) while the other would not. Using these, we can reverse engineer what it is that makes employing the epistemic stance so fitting

to the one, but not the other, and, crucially, what the relevant differences are that make us change the subject from the members to the group. We can turn back to (RQ''). Here it is again:

(RQ''): (i) What would make groups, displaying the appropriate abilities, ever be a successful target of the epistemic stance, and (ii) what should make us change the subject (the relevant entity to be attributed with understanding) from a group's member(s) to the group collectively?

We will answer the two sub-questions (RQ''-i) and (RQ''-ii) in section 3 and 4 respectively, by making use of our two extreme cases and noting their differences. If we have a good argument for what the subject-changing factors are in these extreme cases, we are in a better position to consider more true to life intermediate cases.

3. Macro-systematicity

The question is not whether the disposition of a group would include any abilities (it is uncontroversial that groups can display abilities), but whether those dispositions (abilities included) reveal the group as an epistemic agent. What would make groups, displaying the appropriate abilities, a successful target of the epistemic stance which unites their dispositions as belonging to one epistemic agent? This is the first part of our research question and I shall argue that the answer to it is macro-systematicity (i.e., a higher-level pattern) and elaborate on what that entails.

3.1. Bypassing Micro-Complexity with Macro-Systematicity

The epistemic stance aims to predict or explain. Nothing would allow us to predict or explain happenings in the world more precisely than discerning and utilizing precise laws of nature (through what Dennett calls the physical stance). At a very small level, a micro level, there may be such laws that exploit patterns or systematicities with a high rate of precision. But what they would boast in precision, they lack in swiftness. Try predicting what an expert will do by taking stock of the entire network and signal strength of neurons in her brain and you will appreciate just how impossibly strenuous and time-consuming that would be (even if it were theoretically possible). We would find ourselves with a lack of time and competence to explain or predict her behavior through this level. Because the Expert Planet is isomorphic to just such an expert, we are faced with the same difficulty in predicting or explaining its behavior through the patterns of what its members are doing. In both, we can appreciate the desire to bypass this complexity in some way. But if we wish to bypass the overwhelming complexity from lower-level patterns, then we need an opportunity to do so. That opportunity is macro-systematicity. 'Systematicity' refers to a pattern which a theory can exploit and 'macro' refers to the larger or higher level (relative to another level) where we find this pattern—meaning that we can exploit such a pattern to explain or predict without needing to know any of the lower-level patterns that implement it. The success of the epistemic stance in predicting or explaining individuals entails that there is some systematicity which the stance can

exploit. Systematicity can come in various kinds. It can be exact or, as in the case of the epistemic stance, approximate. While it may not be as precise a predictor as the physical stance, we will see that it makes up in swiftness what it loses in precision.

A pattern which a theory can exploit can be atomistic or holistic (Dennett, 1989; Tollefsen, 2002). The attributions here are holistic, in that their success is to be judged not by how well we can explain or predict the entity's *particular* behaviors as particular attributes, but instead by how well the whole package of attributions fare in explaining or predicting the entity's many behaviors. If the epistemic stance were atomistic, then for every component of the stance there would be an appropriate corresponding behavior (e.g., a belief that *p* corresponding with a particular disposition, such as asserting that *p*). The story, however, is not quite that simple. Ascriptions depend on each other (e.g., how a belief that *p* manifests itself may depend on whether there is a belief that *q*), so it is only by looking at how well sets of ascriptions fare in accounting or predicting groups of behaviors that we can validate those ascriptions. So, the stance does not quite require a 'pattern of behaviour' (Pettit, 2015, p. 100), but a systematicity—as evidenced by the continued success of the stance in predicting or explaining the entity's behavior. For instance, particular beliefs (e.g., that the earth is round) do not require pairing up with particular dispositions (e.g., asserting that the earth is round). It is the package of ascriptions that does the explaining or predicting of the multitude of behaviors.

The ascriptions are holistic in another sense (as opposed to localist) as well: they are spread across their implementation. If the ascriptions were localist, then we could locate each component of a theory in a component of the entity. However, property reductionism must stop somewhere. Brains may store beliefs, but it is unlikely that neurons do and certain that atoms do not. At some point, while decomposing entities into parts, we lose certain concepts because their implementation is spread or distributed *across* the entity we are decomposing. So, the particular attributions we interpret the epistemic agent to have do not need to correspond with physical subcomponents of that entity. Because these systematicities are macro (i.e., of a higher level), we do not need to know any of the lower-level patterns that implement them. The epistemic stance postulates a useful abstraction (the agent) with no dictates on how this must be implemented physically. Groups do not necessarily need to replicate detailed brain-processes before they can be regarded as epistemic agents. Nor do they need to follow any particular joint intention (List & Pettit, 2011), shared plan (Bratman, 2013), we-mode (Tuomela, 2013), or joint commitment (Gilbert, 2013) to bring these macro-systematicities about, as long as they do—which may equally be accomplished by distributing the labor or sub-tasks involved (Hutchins, 1995; Bird, 2014).

The two cases showcase this need for macro-systematicities in the extremes. The composite class does not display any macro-systematicity in the way we would expect from an epistemic agent. The exam, taken as a whole, is too fragmented. We might be able to predict or explain what is going on with individual members and aggregate these explanations to talk about the class, but no benefit is to be gained from seeing the class as one whole. The Expert Planet, on the other hand, does behave, as a whole, with macro-systematicity—namely with the *same* systematicity as the human expert that its brain (and body) is isomorphic to. Hence, the epistemic stance can exploit the same (relevant) macro-systematicities and the epistemic stance would have an equal explanatory and/or predictive power than it would have for an individual.

3.2. *The Shorthand of Reducibility*

If macro-systematicities are what makes an epistemic stance successful, it should follow that there is nothing more to an epistemic agent than the appropriate macro-systematicities. However, there is one point we have not yet considered: the group-level systematicities are not necessarily discernible *only* at the level of the group. If we can give a complete causal explanation couched in terms of cognitive states of individuals (along with physical structures), then we do not need to commit to a new kind of entity or state. This is Rupert's (2011) Simplicity-based Argument. Furthermore, if groups are realized by their members, what characterizes a group or its behavior is always determined by what its members are doing.

In particular, no group agent can form propositional attitudes without these being determined, in one way or another, by certain contributions of its members, and no group agent can act without one or more of its members acting. (List & Pettit, 2011, p. 64)

This is because the group supervenes on its members. Does not that entail that all group-talk is merely a shorthand to talk about its members? In other words, do group properties not reduce to member properties? In the case of individuals, it is easy to see why we do not reduce the epistemic properties (e.g., beliefs) of individuals to the epistemic properties of its parts, because the parts (e.g., neurons) do not invite being ascribed with such properties. But the parts of a group are individuals and individuals are appropriate candidates for epistemic properties. So conceptually it is possible for a group to be a successful target of the epistemic stance *and* have its epistemic properties reducible to the properties of its members. We can easily think of examples where the attributions we make to the group are mere shorthands, for instance when individuals share a purpose (e.g., 'Our class has English at 9:30'), an attribute (e.g., 'Our family isn't particularly bright') or even just a space (e.g., 'Visitors need to leave the premises by 17:00') and we group them according to this shared property. In all these claims we could easily swap the group-label ('class,' 'this family,' and 'visitors') with 'the members of that group,' because the group-talk is a mere shorthand to talk about all of its members. The Composite Class was an extreme example of reducibility because we refer to the whole knowing full well that this is a mere shorthand for a component-story: the composite of its members. When, for instance, we say "these are the exam-answers of the class," we can replace this with "these exam-answers are the composite of the answers of member A, member B,... member Z". As such, the Composite Class is a good example of reducibility. However, it is, as we have shown, a bad example of macro-systematicity. So, before we will be able to answer why we would ever change the subject from the members to the group, we had better use a different case to mark the additional difference; a case where there is reducibility, but also macro-systematicity. Consider the Summative Class:

Summative Class: The answers that the class fills in will be those that are agreed upon by all or most of the members.

This example is constructed in the spirit of the Simple Summative Account (a group believes that p if and only if all or most of its members believe that p - Gilbert, 2013). For our present purposes, let us assume two additional things apply to this class: (i) the members are honest and communicative, and (ii) care is taken so that any decided upon group-answer does not

suffer from any unseemly⁴ inconsistency with its other answers. The singular and (mostly) coherent voice of the Summative Class will give the epistemic stance explanatory power. But employing an epistemic stance towards the group has no benefit over employing an epistemic stance towards (the majority of) its members along with the odd consistency-considerations. Even though we *can* detect group-level macro-systematicities, they are dismissible as superfluous on the ground that we can easily reduce them to the individual-level systematicities. If the purpose of the epistemic stance is to explain or predict, then being able to detect the same properties and make the same predictions (with the same or comparable ease) without going one level higher would make the employment of the epistemic stance towards the group redundant, save as a shorthand or abbreviation. The Summative Class is the shortest hand I can plausibly think of for a group with macro-systematicity. But the hand of group-talk is not always that short, and sizes vary.

4. Emergence

Labelling a set of members as ‘one group’ has the advantage of bypassing the cumbersome talk about individual members with the swiftness of talking about all of them at once, either as a shorthand or because the power of the epistemic stance is uniquely tied to the group-level. Explicating when it is the latter is exactly what is demanded of the second part of our research question (RQ’’-ii). I will argue that the epistemic stance employed towards the group is not only useful, but explanatorily unique to the group-level when the power of the epistemic stance towards the group has no equally powerful counterpart at the member-level, making it irreducible and thus *emergent*.

4.1. The Longhand of Emergence

Emergence is a philosophical term that is often invoked, but hard to define. The general characterization is that “emergent entities (properties or substances) ‘arise’ out of more fundamental entities and yet are ‘novel’ or ‘irreducible’ with respect to them” (O’Connor & Wong, 2015) so that the whole is more than the sum of its parts. What that ‘more’ means, however, is the tricky part. If groups can have ‘novel’ or ‘irreducible’ properties with respect to its members, then it is clear why emergence is so important in why we might change our focus from members to groups. However, it also puts a lot of stress on what we mean by ‘novel’ or ‘irreducible.’ If the characteristics of the groups have to be realized by what their members are doing, then there is always some kind of story at the lower level that accounts for the macro-systematicities at a higher one—thus making it reducible in some way. For authors wielding ‘emergence’ this has created a palpable tension between either trying to deny a reducibility-story (to the extent of making ‘irreducible’ indistinguishable from ‘magical’) or overstressing the importance of finding one (to the extent that any higher-level explanation, no matter how useful or conceptually distinct, would be considered irrelevant if there is also a lower one, no

⁴ While it is true that human individuals are rarely wildly inconsistent (and would be accused of a divided mind if they were), they are not meticulously consistent either. Groups should not be held to a higher standard of consistency than individuals.

matter how complex). With my own characterization of emergence I aim to acknowledge that there is always some relation between the micro and macro, but that the degree of complexity in that relation should make a difference in how we conceive of the macro.

We established that what a person does could be explained or predicted by interpreting her to have beliefs, intentions, and rationality. However, her macro-systematicities are primarily realized by neurons, so there must be *some* elaborate way to talk about beliefs at the neuron level. But, of course, it is not the neurons that have those beliefs. Perhaps we may be able to reduce a person's beliefs to a particular belief-area in her brain, but we cannot locate properties of wholes in their divided parts, ad infinitum. At some level, the property of the whole will be spread across those parts. Furthermore, there is no reason to suppose that patterns at the higher-level can be explained by any systematic grouping (or function of) lower-level patterns. As far as we know, the beliefs of human individuals have no easy mapping-relation to the patterns of its neurons or any selection thereof, making the endeavour of talking about them at the neuron-level as needlessly complicated, unsystematic, and unrewarding as trying to recount the plot of a film using the pixel-level. Some properties will only remain conceptually meaningful at a suitably level of explanation. It is in this sense, and no stronger metaphysical sense, that the epistemic properties of humans are irreducible to neurons or atoms and thus "emergent." That is why I characterize emergence thusly:

Emergence: An entity has emergent properties at a macro-level if those properties have no straightforward mapping-relation to any properties at a level below.

This focuses on the degree of systematicity between two levels of explanation. If this systematicity is high, then we have an easy way to map concepts of the higher-level to those of a lower (namely via that mapping-relation), but if it is not, there is no explanatorily powerful way to talk about the macro-systematicities at the micro-level (because there is no straightforward way to map one onto the other). In short, if reducibility was a shorthand, then emergence is a longhand.

The difference between an emergent and reducible epistemic agent is thus not an operational or functional difference, but a difference of explanatory benefit. We can have two functionally equivalent groups, one which requires an epistemic stance towards the whole (e.g., Expert Planet) and another which does not (e.g., Summative Class). In the latter, the abilities of the group can be systematically pinpointed as abilities of the members. And although an epistemic stance referencing beliefs, aims and tactics may be useful, we can systematically pinpoint how the group's members (and their interaction) feature in this story without losing the conceptual story involving beliefs, aims and tactics. In short, there is a straightforward (and thus explanatorily powerful) way to map one onto the other. In the most straightforward relation, there is a one-on-one correspondence between the group-property and member-properties. Chunking the members into a group saves some labor, but nothing changes conceptually.⁵ In the Expert Planet, however, there is no one-on-one correspondence, nor is there any straightforward way to map those group-concepts on the member-level (any more than there is a way to map individual beliefs to neurons), so the macro-systematicities we detect at the macro-level are conceptually tied to that level. As such, if we want to benefit from the epistemic stance's

⁵ Rather like the relation between machine language and assembly language. (Hofstadter, 1999, p. 290-291)

explanatory power, we can do so only by focusing on the group as a whole. The members simply are not, in any straightforward way, involved in producing the answers on the exam. What the beliefs of the members are with regard to the exam-answers does not (directly) matter. It may indirectly matter in that they might not perform the role of their respective neuron if their opinions or beliefs were otherwise, but it is still in the role of a neuron that they contribute to the larger whole, not as an epistemic agent with their beliefs taken into consideration. Because the China Brain (and thus The Expert Planet) is so unlikely, List and Pettit (2011) decided it best to restrict the conceptualization of group agents to groups of which the attitudes bear some relation to the member attitudes. It bears noting that such a restriction does not only exclude the unlikely cases (like the Expert Planet) from our conceptual scope, but any situation where even one macro-systematicity is more reliant on the actions of the members indirectly (e.g., through self-organization) than the attitudes of the members directly. Therefore, I do not subscribe to this restriction.

The aforementioned characterization fits with all three of aspects of emergence discussed by Theiner (& O'Connor, 2010; Theiner, 2017), namely (a) organizational-dependence, (b) novelty, and (c) autonomy. It is (a) organizational-dependent, because only when the group-properties depend on the *organization* of the members (and not just on an aggregation of their properties), can there be a failure of straightforwardly mapping group-properties onto member-properties. The group level has (b) *novelty*, because if an emergent property is conceptually tied to the higher level, it is new with regard to the previous one. Furthermore, this entails that the group level consequences are not due to the intentions of the members, because if they are, there is a mapping-relation which involves whatever plan the members have to achieve these group level consequences.⁶ Lastly, the group-level is (c) *autonomous* because it is the product of one epistemic stance, detached from the epistemic stances we employ towards its members (since the epistemic stance makes no dictates on implementation). This entails that the group-properties could be multiply realizable. I here depart from Tollefsen (2002), who focuses mainly on (c) multiple realizability to argue for changing the subject, whereas my account, although it entails multiple realizability (otherwise there would be a systematic way to map every macro-systematicity to the lower level), it is still possible that a multiply realizable macro-systematicity has a straightforward mapping relation in each case (e.g., when different sets of members underlie the group-belief), denying the explanatory power as being unique to the macro-level.

In sum, an *individual* has emergent epistemic properties if those properties cannot straightforwardly be mapped to any properties at the neuron/atom level so that we could talk about those properties via neurons. A *group* has emergent epistemic properties (is explanatorily unique) if the stance adopted towards the group has no direct mapping-relation, no shorthand, to any of the stances we would adopt towards its members. There is a difference between using 'group' as shorthand or abbreviation for member-systematicities and using 'group' because member-systematicities are conceptually far removed from the group-systematicities. One may still object to group-talk (as more than a shorthand) by saying that one's interests begin and end

⁶ They may recognize those consequences and be consistently happy with them—which is one way that the members continuing to play their part can be ensured—but that is different from those consequences having been their intention all along.

with human individuals (i.e., members), as opposed to any macro-entity (i.e., the group as group). This is fair enough, as long as one also forgoes talking about the useful macro-systematicities that belong to the group (and not the member) level.

4.2. Degrees & Dimensions of Reducibility

If emergence has to do with how straightforward the mapping-relation is between the micro and macro, then this suggests that there are degrees to emergence. Furthermore, mapping relations can come in various kinds. I believe exploring these degrees and dimensions of reducibility to be a valuable avenue for epistemologists. Summative accounts get contrasted with non-summative ones and critics have attacked non-summative accounts for the way in which their group-story could be reduced to a member-story. (Goldman & Blanchard, 2018) It seems readily acknowledged that such criticisms do not necessarily close the conceptual possibility for group-talk altogether, but why it does not is not clear. If emergence is a key motivator in considering a group as a group rather than as a shorthand for its members, then explicating the manner and degree of emergence would go a long way to explain *why* certain non-summative groups are worthy of their epistemic stance and others are not. With that in mind, I would like to offer up two dimensions (and some degrees within them) that I believe to be at play: (i) selective reductionism, and (ii) transformative reductionism.

The first dimension, (i) selective reductionism, deals with the process that determines which member-property a group-property could reduce to. Here we are starting from the assumption that a group-property can be located in one (or several) of the members and that there is either an explicit or implicit selection procedure that determines which particular or aggregate of member-properties is to stand for the group. There can be varying degrees of complexity to this procedure. The most straightforward selection-procedure would be if there were only one operative member. In such cases, the group-view is just a shorthand to say ‘that particular member’. A case like the Summative Class (or any kind of majority vote) is not far ahead in that it simply selects for the group property that property which is present in the majority or all of its members, along with a consistency requirement. More complex still is when there is an implicit or explicit rule for who is operative in which situation (e.g., let member A answer all the theoretical questions, let B make all the exercises, let C formulate all the reflective questions, and so forth). A little further on, we might find a group where the distribution of properties has extra determining factors in dividing the labour, meaning it has further procedures (explicit or implicit) to decide which property may reduce to which member in which circumstances. We can take this further still and divide each property in sub-components (or sub-tasks) and distribute these among the members. What we can see here is an increasing reliance of the properties of the group on the procedures for selecting them. The more complexity there is to that procedure, the more what characterizes the group is relegated towards the procedure (however that procedure is implemented, be it by rules or by organic dynamics) rather than to the members directly, and the less benefit we would get from reducing the group to its members. Furthermore, in group dynamics, that process of selection could be borne of the interaction between members, which may not admit of any systematic procedure-description to specify a clear mapping-relation.

To make this more intuitive, consider an analogy with composing a text. If I copy-paste *Hamlet* and publish it, Shakespeare is the author of the text, not I. If I compose an Anthology by selecting several texts that make a coherent whole, then an editor credit is certainly called for, but it is still clear that the authorship belongs largely to those whom I selected the texts *from*. If I manage to create a coherent whole by using every other sentence from different existing texts, then there may be a way to reduce every single sentence to its original author, but it is my ingenious selection-process that lies behind the character of the novel as a whole - this despite the fact that I can reduce each sentence in the text to a specific author.

The second dimension, (ii) transformative reductionism, deals with the process that determines how a group-property (or display of ability) could be mapped onto a property of a different *kind* in its members. The smallest amount of transformation is of course no transformation at all, when group properties reduce directly to the same member-properties, as was presupposed in all cases of selective reductionism. The beliefs of the group are provided by the beliefs of the members, the aims of the groups by the aims of the members and so on. A slightly more transformative relation is one where group-properties do not reduce to the same member-properties, but very similar ones nonetheless. An example of this is *beliefs* of the group reducing to the *commitments* of (the majority of) their members. As Gilbert (2013) has pointed out, members can (jointly) commit to something they don't believe, in which case the group can be said to believe something that none of its members do. But is this sufficient for an emergent property? In a weak sense, yes, but there is still only a short distance to cross in explaining the group-level story with a member-level story: namely, the commitments of each of the members to act in a way that reflects that belief along with their expressed readiness to do so (on pain of rebuke) due to the commitment. (Gilbert, 2013) The fact that Gilbert is able to convey what constitutes a group belief in such individualistic terms betrays the ease with which we can map the group properties to a set of individual properties. Although group properties are a useful shorthand and help not to conflate which properties function as those of the group with those of the members, we can still map one onto a set of the other. The case of the *Averaging group* (the group-belief is the average of the member's beliefs) is another example of a short distance between the member- and group-story. Taking the average of member-beliefs can result in a 'novel' belief not attributable to any of the members (e.g., when asked for the solution of 3+2, if A answers 4 and B answers 6, then the group settles on 5), but behind its novelty is a transformation-relation that is too small to feel the need to change the subject. We can still easily use group-talk as a shorthand for what the members believe—by considering it as an average of the member-answers. The more complexity we add to that procedure, the more what characterizes the group is relegated towards the procedure rather than the members and the more difficult to talk about the macro-systematicity in groups via those of its members. Furthermore, if the procedure is due to continuous interaction between members, it may not admit of any systematic procedure-description to specify a clear mapping-relation.

The same authoring text metaphor can be used here. If I copy-paste *Hamlet* and print it, I am not the author of the text. Nor would I be if I added a small transformation to each word (e.g., changing the font). Translating the text into another language would earn me a Translator credit, but it would not fundamentally change the author of the text. An adaptation, on the

other hand, *would* start doing that (and move Shakespeare into a less glamorous credit of “Inspired by the work of”). The more the work gets transformed, the more credit goes to whatever transformation-process was involved.

In both dimensions, what creates the emergence is the increasing complexity in the mapping-relation. In selective reductionism, it is the way in which it is decided which member-property will operate as the group property. In transformative reductionism, it is the way in which member-properties are molded to operate for the group property. These dimensions intermingle quite a bit. Much like the two extremes earlier, I think the extremes along these dimensions are interesting because they are conceptually clearer, but I would be quite surprised to, for instance, see a group with no transformation-relations, but a high selection-relation for similar reasons that I would be surprised to see texts composed of other people’s sentences.⁷ Additionally, nothing stands in the way of the various group-properties having different kinds of relations to its member-properties.

Conclusion

Could groups, collectively, understand? We have elucidated what that question means by first specifying what it means for an entity to understand, namely: possessing the appropriate abilities. Groups can certainly display the relevant abilities, but this does not automatically entail that groups, rather than their members, understand anything. Labelling a set of members as ‘one group’ has the advantage of bypassing the need to cumbersomely talk about individual members with the swiftness of talking about all of them at once. But is group-talk ever more than just a shorthand or a metaphor? Conversely, is epistemic agenthood ever something that could or should be attributed to a group (as opposed to its members)? To answer this, we have explicated what it means to be an epistemic agent: being a successful target of the epistemic stance. The epistemic stance creates an instrumental abstraction by attributing epistemic properties (i.e., beliefs, epistemic aims, problem-solving tactics) to an entity to explain or predict its behaviour. The explanatory or predictive power of the stance is what validates any entity (such as a group) as an epistemic agent. Macro-systematicity (a higher level pattern which a theory can exploit) is what makes the epistemic stance’s abstraction explanatory powerful (regardless of how that macro-systematicity is realized) and emergence (the lack of a straightforward relation between the micro and the macro) is what makes the power of the epistemic stance unique to a particular level of explanation (because the relation between macro and the micro is so complex that the concepts at one level are no longer appropriate to describe the other). But there are different degrees and dimensions to this relation (two were suggested in this paper: the complexity in the selection and transformation procedure) and these require further development. While I have not answered whether there do exist good candidates of group epistemic agents, I hope to have shed light on the conceptual space involved in substantiating such an answer.

⁷ Except as an exercise precisely with this methodology in mind.

References

- Bird, A. (2014). When is There a Group that Knows? Distributed Cognition, Scientific Knowledge, and the Social Epistemic Subject. In J. Lackey (Ed.), *Essays in Collective Epistemology* (pp. 42-63). Oxford University Press. Could not find a publisher location in this book chapter reference.
- Block, N. (1978). Troubles with functionalism. *Minnesota Studies in The Philosophy of Science*, 9, 261–325.
- Bratman, M. E. (2013). *Shared agency: A planning theory of acting together*. Oxford University Press. Could not find a publisher location in this book reference.
- Dennett, D. C. (1989). *The intentional stance*. Cambridge, MA: MIT Press.
- Dennett, D. C. (2009). Intentional systems theory. In A. Beckermann, B. P. McLaughlin, & S. Walter (Eds.), *The Oxford Handbook of Philosophy of Mind* (pp. 339–349). Oxford: Oxford University Press.
- Delarivière, S., & Van Kerkhove, B. (2017). The “Artificial Mathematician” Objection: Exploring the (Im)possibility of Automating Mathematical Understanding. In B. Sriraman (Ed.), *Humanizing Mathematics and its Philosophy* (pp. 173–198). Basel: Birkhäuser. <https://doi.org/10.1007/978-3-319-61231-7>
- De Regt, H. W., & Dieks, D. (2005). A contextual approach to scientific understanding. *Synthese*, 144(1), 137–170. <https://doi.org/10.1007/s11229-005-5000-4>
- Gilbert, M. (2013). *Joint commitment: How we make the social world*. Oxford: Oxford University Press.
- Goldman, Al. & Blanchard, T. (2018). Social Epistemology. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Summer 2018 Edition)*. Retrieved from: <https://plato.stanford.edu/archives/sum2018/entries/epistemology-social/>
- Grimm, S. R. (2014). Understanding as knowledge of causes. In A. Fairweather (Ed.), *Virtue epistemology naturalized* (pp. 329–345). Dordrecht: Springer.
- Hills, A. (2016). Understanding Why. *Noûs*, 50(4), 661–688. <https://doi.org/10.1111/nous.12092>
- Hofstadter, D. R. (1999). *Gödel, Escher, Bach: An Eternal Golden Braid*. New York: Basic Books.
- Hutchins, E. (1995). *Cognition in the Wild*. Cambridge: MIT press.
- Khalifa, K. (2013). Understanding, grasping and luck. *Episteme (Edinburgh)*, 10(1), 1–17. <https://doi.org/10.1017/epi.2013.6>
- Kvanvig, J. L. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge: Cambridge University Press.
- List, C. & Pettit, P. (2011). *Group Agency. The Possibility, Design and Status of Corporate Agents*. Oxford: Oxford University Press.
- O’Connor, T. & Wong, H. Y. (2015). Emergent Properties. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Summer 2015 Edition)*. Retrieved from: <https://plato.stanford.edu/entries/properties-emergent/>

- Pettit, P. (2014). How to Tell if a Group is an Agent. In J. Lackey (Ed.), *Essays in Collective Epistemology* (p.97–121). Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199665792.001.0001>
- Rupert, R. D. (2011). Empirical arguments for group minds: A critical appraisal. *Philosophy Compass*, 6(9), 630–639. <https://doi.org/10.1111/j.1747-9991.2011.00420.x>
- Ryle, G. (2000). *The Concept of Mind*. London: Penguin.
- Van Camp, W. (2014). Explaining understanding (or understanding explanation). *European Journal for Philosophy of Science*, 4(1), 95–114. <https://doi.org/10.1007/s13194-013-0077-y>
- Wilkenfeld, D. A. (2013). Understanding as representation manipulability. *Synthese*, 190(6), 997–1016. <https://doi.org/10.1007/s11229-011-0055-x>
- Theiner, G. (2017). Groups as Distributed Cognitive Systems. In Ludwig, K. & Jankovic, Marija (Eds.), *The Routledge Handbook of Collective Intentionality* (pp. 233-248). New York, Routledge.
- Tollefsen, D. P. (2002). Collective intentionality and the social sciences. *Philosophy of the Social Sciences*, 32(1), 25–50. <https://doi.org/10.1177/004839310203200102>
- Tollefsen, D. P. (2015). *Groups as Agents*. Cambridge: John Wiley and Sons.
- Tuomela, R. (2013). *Social ontology: Collective intentionality and group agents*. Oxford University Press.
- Trout, J. D. (2002). Scientific explanation and the sense of understanding. *Philosophy of Science*, 69(2), 212–233. <https://doi.org/10.1086/341050>
- Ylikoski, Petri. (2009). The Illusion of Depth of Understanding in Science. In H. de Regt, S. Leonelli and K. Eigner (Eds.), *Scientific Understanding: Philosophical Perspectives* (pp. 100–119). Pittsburgh: University of Pittsburgh Press.
- Zagzebski, L. (2001). Recovering understanding. Knowledge, truth, and duty. *Essays on epistemic justification, responsibility, and virtue*, 2001, 235-252.

The editorial and publishing process of this publication has been financed by the Ministry of Science and Higher Education from the funds for the dissemination of research (DUN) within the framework of publishing activity, contract no. 711/P-DUN/2019, period of implementation: the years 2019–2020.