

# Philosophy of the Social Sciences

<http://pos.sagepub.com/>

---

## **Collective Actors without Collective Minds: An Inferentialist Approach**

Javier González de Prado Salas and Jesús Zamora-Bonilla

*Philosophy of the Social Sciences* published online 20 February 2014

DOI: 10.1177/0048393113520397

The online version of this article can be found at:

<http://pos.sagepub.com/content/early/2014/02/12/0048393113520397>

---

Published by:



<http://www.sagepublications.com>

**Additional services and information for *Philosophy of the Social Sciences* can be found at:**

**Email Alerts:** <http://pos.sagepub.com/cgi/alerts>

**Subscriptions:** <http://pos.sagepub.com/subscriptions>

**Reprints:** <http://www.sagepub.com/journalsReprints.nav>

**Permissions:** <http://www.sagepub.com/journalsPermissions.nav>

>> [OnlineFirst Version of Record](#) - Feb 20, 2014

[What is This?](#)

# Collective Actors without Collective Minds: An Inferentialist Approach

Philosophy of the Social Sciences  
1–23

© The Author(s) 2014

Reprints and permissions:

sagepub.com/journalsPermissions.nav

DOI: 10.1177/0048393113520397

pos.sagepub.com



Javier González de Prado Salas<sup>1,2</sup> and Jesús Zamora-Bonilla<sup>1</sup>

## Abstract

We present an inferentialist account of collective rationality and intentionality, according to which beliefs and other intentional states are understood in terms of the normative statuses attributed to, and undertaken by, the participants of a discursive practice—namely, their discursive or practical commitments and entitlements. Although these statuses are instituted by the performances and attitudes of the agents, they are not identified with any physical or psychological entity, process or relation. Therefore, we argue that inferentialism allows us to talk of collective intentionality and agency without needing to posit the existence of any sort of collective psychology or mind.

## Keywords

normativity, inferentialism, rationality, deontic status, belief, intention, commitment, entitlement, collective intentionality, psychological states

## 1. The Problem of Collective Agency

One of the most important problems in the philosophy of social sciences is the nature of collective entities, properties, and actions, and its relation to

---

Received 7 May 2013

<sup>1</sup>Universidad Nacional de Educación a Distancia, Madrid, Spain

<sup>2</sup>University of Southampton, Southampton, UK

### Corresponding Author:

Javier González de Prado Salas, Departamento de Lógica, Historia y Filosofía de la Ciencia, Facultad de Filosofía, Universidad Nacional de Educación a Distancia, Paseo Senda del Rey 7, Madrid 28040, Spain.

Email: jgonzalezdeprado@gmail.com

those of individuals.<sup>1</sup> In the tradition of analytical philosophy, empiricism, and naturalism, the tendency has been to adopt some or other variety of ontological or methodological individualism, according to which only individual human beings can be agents, that is, can have intentional states and govern their behavior out of these, whereas collective events would just consist in the (more or less complex, but mechanical in the end) aggregation and interconnection of individual activities and stances. Other authors, typically influenced by traditions in continental philosophy, have instead argued that collective actors, intentions, and actions are emergent realities having a peculiar nature by themselves, irreducible to those of single human beings, in a similar fashion, after all, as the individual mind emerges as an autonomous and specific causal agent out of the biological and chemical entities (e.g., cells) of our organisms.

In this article, we want to offer an explanation of collective actors that respects the empiricist-naturalist claim that there are no such a thing as collective “minds,” that is, mental states that are not the states of some or other individual mind (with all the kinds of mutual dependencies that science may discover between these), but that allows for the existence of a specific type of rational agency that can directly be attributed to collective entities. The kernel of our argument will be the thesis that what transforms a series of “natural” events (e.g., neural or psychological states, describable in principle by the natural sciences) into an intentional state or an example of *rational* agency is the *normative import* attributed to those events, and that the objects of these attributions are not constrained by the kind of “psychological reality” those objects may have, or at least, the constraints are not so tight that only individual mental states are the types of events to which the normative import can be attributed.

Our argument is based on the inferentialist theory of rationality and normativity mainly developed by Robert Brandom (1994), and hence, we will start by offering a summary of that theory in Section 2. In Section 3, we will focus on how intentional states—such as beliefs and intentions—are accounted for in inferentialism: the fundamental claim here is that attributions of intentional states are to be understood as attributions of normative statuses, rather than as attributions of a particular mental or physical state. Then, in Section 4, we will study collective intentionality from this inferentialist perspective. Finally, in Section 5, we will contrast our approach with other contributions to the debate on collective intentionality, and we will assess the comparative advantages of our proposal.

---

<sup>1</sup>For two recent surveys of approaches to this and related questions, see Bouvier (2011) and Kaufmann (2011).

## 2. A Brief Sketch of Inferentialism

There is an influential philosophical view that characterizes rational agency as purposive behavior, involving representational capacities. Such view is particularly influential in the debate on collective intentionality. For instance, List and Pettit (2011) argue that, to be a rational agent, a system must be able to have representational states that describe how the world is, and also motivational states that specify how the system requires the world to be; furthermore, the system must be able to process these states so as to produce actions that aim at the satisfaction of the motivating specifications.

Most standard accounts of rationality tend to go along these representationalist lines. However, in this article, we will take an alternative approach—which, we think, offers a more promising angle on the study of collective intentionality. In particular, we will consider inferentialist accounts of rationality and intentional content. One of the more developed presentations of inferentialism is the one found in the work of Brandom (1994)—who draws heavily from the pioneering contributions of Frege, Wittgenstein, Sellars, and Dummett. This is the version of inferentialism we will focus on here. In this section, we will try to present some of its main ideas, which will later be applied to the analysis of collective intentionality. It should be clear that we do not intend to offer here reasons in favor of an inferentialist approach, but rather to examine whether such an approach is illuminating for the study of collective intentionality.

Brandom's inferentialism holds that the fundamental mark of rationality is the capacity to engage in (social) games of "giving and asking for reasons." Thus, rational agents, to whom contentful intentional states may be ascribed, would be those who can take part in such practices of giving and asking for reasons. An agent counts as *rational* if she can be held *responsible* for her actions and beliefs—that is, if she can be demanded to provide reasons for them. Rational agency, as understood here, is therefore closely linked to normativity, as the justificatory games of giving and asking for reasons are intrinsically normative (as will be explained below).

The content of intentional states or performances is determined—in this inferentialist framework—by their inferential role in these justificatory practices: that is, the content of a claim *p* would be given by an account of which other claims provide reasons to endorse *p*, and for the endorsement of which further claims *p* may count as a reason (or as part of a reason). Mastering such inferential relations is what makes you a rational and discursive being.

Accordingly, Brandom (1994, 93-94) takes inference<sup>2</sup> as his semantic primitive, in contrast to the main tradition in modern philosophy, which has usually chosen representational notions (such as reference, picturing, satisfaction, isomorphism, or designation) as their primitives. Semantic content is characterized in terms of inferential role. In this way, a given state or performance can be said to express certain content  $p$  when the role of such state or performance in the game of giving and asking for reasons mirrors the inferential relations in which the claim  $p$  is involved. Thus, if the claim  $p$  entails the claim  $q$ , a state or performance expressing the claim  $p$  may count as offering reasons for a state or performance expressing the claim  $q$ .<sup>3</sup>

Assertions would be the kind of speech acts whose contents can play the role both of premises and conclusions in inferences—so they can be offered as, and can stand in need of, reasons. Hence, semantic content will be associated primarily with assertions, which are the fundamental move in discursive practices—whereas the significance of other speech acts would derive from that of related assertions (Brandom 1994, 171-72).<sup>4</sup>

Brandom, as we have seen, grounds his semantic theory on an account of the social practices of giving and asking for reasons. Furthermore, he tries to analyze such social practices in purely pragmatic—nonsemantic—terms. He is following here a self-imposed pragmatist constraint, according to which semantic notions should be elucidated in terms of use or pragmatic

---

<sup>2</sup>The same is true regarding incompatibility, though whether this notion and the notion of entailment are mutually reducible to each other is a question not directly related to our discussion.

<sup>3</sup>Strictly speaking, reasons for a given state/performance are not given by other states/performances, but by the claims they express. However, the way in which reasons are offered is through a further state/performance which expresses the relevant reason-giving claim (typically, asserting such reason-giving claim). So, states/performances can be said to provide reasons for other states/performances, in a derivative sense. We say "may count" because, to provide reasons, the first state/performance has to be well grounded itself.

<sup>4</sup>Kuklla and Lance (2009) reject this primacy, arguing that other speech acts, like what they call "observatives" or "prescriptives" are different in kind from "declaratives" and independent of them. Although we are sympathetic to their view, we prefer not to complicate our argument with that discussion.

significance<sup>5</sup>: semantics must answer to pragmatics (Brandom 1994, 144–47; see also MacFarlane 2010). If Brandom is successful, he will be able to describe in nonsemantical (or intentional) terms the structure of the practices that count as expressing certain semantic content.<sup>6</sup> However, this does not automatically imply that Brandom’s project is reductionist, at least not reductionist in a straightforward naturalist sense. The structure of the practices that give rise to the expression of semantic content is, according to Brandom, intrinsically normative. That is, we need to resort to normative vocabulary to give a suitable characterization of such practices. For, in a practice that includes giving and asking for reasons, participants do not only produce performances but also assess the performances of others: in certain circumstances, some performances are deemed to be appropriate and some others inappropriate, and a given performance would be regarded as having certain appropriate consequences.

“Reason” might be considered a notion suspect of intentional leanings, and therefore it is not the best candidate for the job of describing semantic practices in purely pragmatic terms. Brandom prefers to construct his pragmatic analysis with other normative terms whose pure pragmatic pedigree is less controversial. The basic normative notion of appropriateness is, however, not fine grained enough, so Brandom resorts to the more refined deontic concepts of commitment and entitlement (Brandom 1994, 159–61).

The notions of commitment and entitlement are related to those of obligation and permission. Somebody is committed to doing A if not doing it is inappropriate, whereas being entitled to doing B means that its performance is not inappropriate. Incompatibility is, in turn, explained as a commitment that precludes entitlement to some other commitment.<sup>7</sup> Assertoric commitments are not isolated, but mutually connected by *inferential norms*. When I assert something, I become committed to that claim, and also to further claims that are entailed by that one (what further commitments follow from

---

<sup>5</sup>A complementary pragmatist constraint imposes that semantic notions should only be introduced with the aim of contributing to the explanation of the pragmatic significance of performances (Brandom 1994, 144).

<sup>6</sup>This is what he has called “elaborating-explicating”—LX—relation between two vocabularies, one of them being capable of describing a practice that is sufficient to master the second (cf. Brandom 2008, 47).

<sup>7</sup>The notion of incompatible or incoherent beliefs—which poses difficulties for causal accounts—is not problematic here: having incoherent beliefs is just doing something inappropriate, but not unintelligible (Brandom 1994, 195–96).

the commitments endorsed by the speaker will depend on the inferential norms active in the practice). In the same way, an entitled assertion provides entitlement for other assertions (which other assertions, again, will be determined by the inferential articulation of the linguistic game). What is a reason for what—that is, what provides entitlement for what—is determined, thus, by the inferential norms of the practice. Taking part in the game of giving and asking for reasons would amount to being subject to such norms.

Brandom calls the set of deontic statuses of a participant as her “deontic score.” The content of a claim would be given by the commitments and entitlements related to its assertion<sup>8</sup>—that is, by how its assertion alters the deontic score of the speaker. The participants of a discursive practice keep track of the evolution of the score of each speaker, in an activity called “scorekeeping” (Brandom 1994, 181-87). The role that these statuses play in discursive practices is accounted for by reference to the attitudes adopted by the participants when keeping score: being committed to something is understood in terms of *being taken* to be committed to it (Brandom 1994, 133-34). The fundamental attitudes are those of *attributing* and *acknowledging* a commitment or entitlement (Brandom 1994, 165-66). For instance, when A attributes to B the entitlement to doing X, A will treat B’s performance of X as appropriate. Furthermore, a participant *undertakes* a certain commitment if she does something in virtue of which it is appropriate for others to attribute such commitment to her. On the other hand, *acknowledging* some commitment/entitlement consists in endorsing it oneself (Brandom 1994, 196, 269-71, 596)—and, of course, acknowledging a commitment implies (at least under normal circumstances) undertaking it, and also attributing it to oneself.

Finally, the attitudes of users are revealed through their sanctioning behavior: if my performances are taken to be incorrect by other participants, they will be disposed to sanction me correspondingly. Some of these sanctions could, in principle, be describable in terms of nonnormative behavior—paradigmatically, reinforcement and punishment—but Brandom (1994,

---

<sup>8</sup>It should also be pointed out that assertoric commitments (the sort of commitments undertaking when making a claim) are not restricted in their scope: that is, any possible hearer may hold me accountable for my assertion (and sanction me accordingly). This universality of assertoric commitments can be contrasted with the narrow focus of the kind of commitments acquired in promises, which are specifically targeted to some individual(s)—the promisee(s). See Wanderer (2008).

34-46) also wants to leave room for “merely” normative sanctions, which would consist in alterations of the deontic score—of the commitments and entitlements attributed to the speaker or to others.

It may seem that discursive practices with the structure shown above are mere intralinguistic games with no contact with the world. However, Brandom (1994, ch. 4) admits noninferential inputs and outputs to the discursive game of giving and asking for reasons, which would make it possible for participants to engage with the “external” world through discursive practices. Linguistic moves are not the only thing that can provide—and stand in need of—reasons. Some nonlinguistic circumstances may offer reasons for a discursive move and, likewise, some linguistic moves provide reasons for nonlinguistic outputs (i.e., actions).

The typical case of noninferential input is observational reports. The perception of a red table—by a reliable observer—would be an appropriate circumstance for reporting “There is a red table.” Even if that report is elicited, and *prima facie* entitled, by a noninferential circumstance of application, it is integrated in the structure of inferential relations: the report would commit and entitle us to further inferential moves (for instance, it would commit us to the claim that there is a colored piece of furniture, or to provide evidence for our claim).

Noninferential outputs, on their part, are related to practical reasoning: the performance of certain actions can be included among the appropriate consequences of application of inferential moves. Some commitments/entitlements are commitments/entitlements to the performance of actions—which are not themselves inferential, but are connected to the inferential web in the sense that we can give and ask for reasons for them. For example, the action of opening the umbrella would be licensed by an entitled commitment to the claim that it is raining.<sup>9</sup> A piece of practical reasoning, then, would be an inferential pattern that connects doxastic commitments with the performance of an action (or with some practical commitment/entitlement).

---

<sup>9</sup>What actions are licensed—or required—by certain doxastic commitments depends, in general, on the goals of the agent—and on the social norms she is subject to. Therefore, if I am hungry, the fact that there is food in the fridge will give me reasons to go to open the fridge. In the same way, the existence of a rule that demands to stop in front of red traffic lights gives me reasons to stop my car in front of a red light. We will not go into the details of Brandom’s (1994, ch. 4) account of these issues.

### 3. Inferentialism and Intentional States

In the previous section, we have presented a brief summary of Brandom's pragmatic analysis of semantic-discursive practices—that is, practices that manage to express intentional content. It is remarkable that the sort of notions that traditionally play a central role in explanations of intentionality—such as “belief” or “intention”—do not figure in this analysis. What is the place of these intentional notions in Brandom's inferentialist framework? As the debates in collective intentionality are usually framed in terms of these traditional intentional notions, this question must be addressed before engaging in those debates from an inferentialist perspective. This is what we will do in this section.

According to Brandom's pragmatic account of discursive practices, the paradigmatic discursive move through which one acquires a commitment to some claim is the assertion of such a claim. As we have seen, when a speaker asserts a claim, she typically becomes committed to that claim (and to its justification) and to all the other claims that follow from it. That is why it can be said that the speaker's assertion expresses that contentful claim.

The role played by assertion in public speech has its psychological counterpart in the attitude of believing<sup>10</sup>. In Brandom's inferentialism, the acquisition of the belief that *p* is understood as the undertaking of a commitment to the claim that *p*. Thus, attributing a belief to some speaker amounts to attributing to that speaker some doxastic commitment—that is, a commitment to some claim. As asserting something implies acknowledging a doxastic commitment, beliefs will be expressed by assertions. In this way, the notions of assertion and belief are intrinsically linked in inferentialism.

Intentions are also explained in terms of commitments, although not doxastic ones. Intentions are conceived as practical commitments, this is to say, commitments to the performance of an action. So, an agent forms an intention when she acknowledges a practical commitment to the performance of an action.<sup>11</sup> An agent will have the intention to open the door when she acknowledges a commitment to opening the door.

---

<sup>10</sup>We will use the term ‘believe’ or ‘believing’ to refer to the mental state or attitude of believing, while the term ‘belief’ would refer to the content of such mental states. Thus, a belief is in general some proposition or claim.

<sup>11</sup>The performance of an action would count as acknowledging a practical commitment to such performance—it would be an “intention in action.”

The traditional intentional states of believing and intending,<sup>12</sup> therefore, are accounted for by appealing to the attribution and acquisition of commitments and entitlements. Notice that commitments and entitlements are deontic normative statuses and, as such, should be distinguished from psychological or physical states, processes, or events. Thus, when we attribute some intentional state to an agent, we are not attributing any sort of such physical or psychological events or states<sup>13</sup>—but rather, as we have just pointed out, deontic statuses. In principle, there is no specific constraint as to what sort of events or states may count in a given practice as the undertaking of a commitment or entitlement. The only requisite is that the agent or system to which these statuses are attributed has the capacity to behave so as to be counted as following the inferential rules that help define its deontic score (of course, this does not mean that, to count as subject to a norm, an agent must always act in accordance with the norm, but rather that her performances are susceptible of being evaluated in relation to such norm and that they are minimally responsive to that evaluation).

At first sight, the identification of our ascriptions of intentional states with attributions of deontic statuses may seem problematic. The reason for this is that the point of interpreting agents by means of intentional states is to describe, explain, and predict their behavior and mental life. But, if our psychological interpretations actually consist of attributions of deontic statuses, it would seem that, at best, we would be describing what agents *ought to do*, not what they *will do*. Deontic statuses would not have direct causal effects, so it is not clear how they will be able to explain the psychological causal processes that give rise to the behavior of agents.

However, we have to take into account that, even if commitments and entitlements are not to be equated with any physical or psychological entity or event, they are not completely unrelated to the behavior and psychology of

---

<sup>12</sup>There would be other (related) inferentialist stories for the rest of the typical intentional attitudes, such as desiring or imagining. See Brandom (1994, 2008).

<sup>13</sup>Of course, we can have a deflationist view of psychology, according to which a “psychological” or “mental” state would just be what is attributed when we attribute, for instance, a belief—so that being attributed a doxastic commitment may count as being in certain “mental” state. Actually, this is one of the conclusions that might be drawn from our arguments in this article. However, we will use the term “psychological” to refer to some more substantial notion, which would require the instantiation of the state or process in question in a nervous system or some functional equivalent.

agents. According to Brandom's pragmatics, deontic statuses are introduced in discursive practices only through the attitudes of agents toward such statuses. Remember that the status of having some commitment was understood in terms of being taken to have such a commitment. Thus, commitments and entitlements only take part in the practice as something that is attributed or acknowledged (Brandom 1994, 133-34). And the normative attitudes of agents are related to the actual behavior and psychology of agents—we have even seen that they may be understood in terms of dispositions and sanctioning behavior. In particular, many of the attitudes of agents are closely connected with their practical performances; for instance, the paradigmatic way of acknowledging a doxastic commitment is by *making* an assertion, while practical commitments are typically acknowledged by *performing* an action. So, it is through the normative attitudes of agents that normative statuses (commitments and entitlements) may have any causal role in the practice (Brandom 1994, 626).

However, the fact that—to use Brandom's (1994, 133) words—normative statuses are “instituted” by the attitudes of the participants does not mean that the relation between attitudes and statuses is always direct. For instance, there can be commitments attributable to an agent that, nevertheless, are not acknowledged by her. In that case, the agent would have undertaken a commitment but this undertaking would not be reflected in her psychology. This situation is possible because when an agent undertakes a commitment to some claim, she also becomes committed to those other claims that follow from it—even if she fails to realize that she has acquired all these consequential commitments. Therefore, a commitment can be undertaken by an agent (i.e., it is attributable to her), but not acknowledged by her (at least, according to the way Brandom characterizes the notions of undertaking and acknowledging a commitment).

This distinction between acknowledged and unacknowledged—but still undertaken—commitments yields two possible characterizations of the notion of belief, which are not always coincident (Brandom 1994, 194-96). On the one hand, when we specify the beliefs of an agent, we may be taking into account the commitments she acknowledges, what she *thinks*. These descriptive specifications appeal to attitudes, but not to statuses, and can be explained by means of dispositions, regularities, and causally efficacious psychological mechanisms. On the other hand, our ascriptions of beliefs may be concerned about the commitments the agent has actually undertaken, even if she does not necessarily acknowledge them—that is, about what she *ought to think* (given what she endorses). Brandom does not regard this ambiguity

in the notion of belief as a problem for his theory, but rather as a real feature of our ascriptions of intentional attitudes.<sup>14</sup>

Furthermore, Brandom wants to preserve a strong notion of objectivity, which allows for the possibility of an agent actually undertaking a commitment without anybody in the community attributing it—or being disposed to attribute it. That is, there may be statuses that would outrun the attitudes of all the speakers in the community—commitments that would be actually undertaken even if nobody acknowledges or attributes them. This objectivistic requisite appears to be in tension with the idea that statuses are instituted by attitudes<sup>15</sup>—and also seems to pose difficulties for the naturalization of inferentialism. However, the discussion of these issues lies beyond the scope of our aims in this article.

So, to sum up, besides the distinction between attributing a commitment to some other agent and undertaking it oneself, we can also distinguish between *acknowledged* and *implicit* (i.e., undertaken but not necessarily acknowledged) statuses. The acknowledgment or attribution of a status corresponds to real mental events, and in this character, they represent the only way in which we can, somehow figuratively, assert that normative statuses can have some *causal* efficacy (figuratively because *it is not the status itself that has causal powers*, but the psychological states that constitute its acknowledgment and its causal links with the rest of the world and society; take into account that the status is not identical to the acknowledgment itself, but to *what is acknowledged* by it). Implicit statuses are those that, given how the world is and what inferential rules are accepted by the subjects, *follow* from the actually attributed or undertaken statuses, but may have not been in fact attributed or *acknowledged*.

The proposal we are defending is that attributions of intentional states are attributions of normative statuses and that, therefore, to characterize intentional agency, one should look at the normative structure of the practice couching such agency—rather than to whatever specific physical or psychological processes underpin that practice. We are not committed to any

---

<sup>14</sup>Actually, Brandom (1994, 196) does not propose to analyze the notion of belief in terms of commitments and entitlements, but rather to abandon the ambiguous idea of belief in favor of these more precise normative concepts. We are more favorable to distinguish beliefs as normative statuses from beliefs as (say) psychologically real states of our neural system, as long as the two can often go in different directions.

<sup>15</sup>For criticism and analysis of this tension, see, for example, Shapiro (2004), Wanderer (2008, 78-92), Laurier (2005), Grönert (2005), Hattiangadi (2007), Rödl (2010), and Rosen (1997).

particular view of these underlying physical and psychological processes. The only thing that is required from such processes—the “hardware” of the practice, as we will put it in the following sections—is that they give rise to agents who can be interpreted as taking part in a normative practice of giving and asking for reasons. Actually, it could well be that there is not a single “hardware” capable of underpinning this sort of normative skills.<sup>16</sup> At any rate, we do not identify the acquisition of intentional states with any of these underlying physical or psychological processes, but rather with certain normative moves in a suitable practice of giving and asking for reasons. What is essential for rational agents is that they are able to engage in such normative practices, and this may be possible for entities that do not share the underlying physical or psychological constitution of typical *individual* human agents—which includes, among other things, an individual brain and certain kinds of psychological processes. In the next section, we will exploit this possibility to offer an account of collective intentionality—that is, an account of attributions of intentional states to entities like companies or governments, which do not have a (single) brain.

#### 4. The Ontology of Commitments: The Collective Case

Intentional states, such as beliefs or intentions, are often attributed to collectives (corporations, companies, nations). For instance, a certain company can be said to believe that the economy will improve next year. However, it is not altogether clear how we can make this sort of attributions to collective entities, which do not have—at least in an intuitive sense—a brain or a mind. Can we make sense of the idea of collective agents without being forced to accept some mysterious form of collective mind or psychology? Or are attributions of collective intentional states just a relaxed way of talking? In this section,

---

<sup>16</sup>It could be argued that, in fact, only agents with certain physical and psychological characteristics will be able to engage in such normative practices. It could even be that there is a specific neurological or mental state corresponding to each state susceptible of counting as the undertaking of a discursive or practical commitment. But this claim needs to be supported by further arguments—and, if the proposal we are presenting here is on the right tracks, there is no guarantee that this sort of investigation will be successful, because the study of isolated individual brains and mental life does not take into consideration the complex social contexts that, in the view we are presenting, form the background of normative discursive practices. A more promising alternative would be to study what sort of creatures—engaging in what sort of interactions—may give rise to (what can be interpreted as) social normative practices.

we will consider how the inferentialist approach to intentionality that we have presented above may help to illuminate this problem.

We have seen that, regarding the ontology of commitments at the level of individuals, it is possible to differentiate two kinds of entities: psychological states and normative statuses, of which only the former (some of which consist in the acknowledgment or attribution of some of the second) have real causal power. What normative statuses the agents actually have will depend on the inferential norms endorsed by specific groups of interacting people. The question whether collective entities can be rational agents, in the sense of having deontic statuses and having the capacity of changing some statuses by their own actions, has from this point of view an almost trivial answer: it depends on the inferential norms that are accepted by people—that is, it depends on the structure of the normative practices of that community.

If some group of people accept the norm that pushing a ball beyond a particular mark of chalk in some particular way entails that one football team increases in one unit its score in a match, then football teams will have deontic scores consisting in such and such numbers, depending on how many times a ball has passed the line in the appropriate way. It is obvious from this example that many things entail the institution or the change in the normative score of collective entities: sport teams, firms, organizations, governments, and so on. These entities have obviously deontic statuses, including those allowing us to interpret some events in the world as *actions* of those entities (winning a match, declaring a war, distributing dividends, buying a skyscraper, etc.). And some of these collective actions may, in their turn, have the character of the collective entity *endorsing* some norms and hence creating or attributing other commitments to itself or to other collective or individual entities. In this sense, there seems to be absolutely no ontological or epistemological limits for attributing deontic statuses, and hence the capacity of rational action, to collective subjects: it will just depend on what inferential norms the participants of the practice in question actually endorse. What is the reason, hence, why there has been so strong a dispute about the possibility of “collective commitments” and the like?<sup>17</sup>

---

<sup>17</sup>In the last years, most authors have tended to accept, in some sense or another, the existence of collective intentions, attitudes, commitments, and the like, though there is still a big gap between those that are happy to admit collective psychological states as *sui generis* entities (e.g., Schmid 2009), and those that attempt to reduce collective states of mind to the complex interaction of individual minds (adding, in general, an individual disposition to enter into what they call a “we-attitude”; for example, Searle 1995, Tuomela 1995). As we shall see immediately, our approach does not coincide with any of these.

Our impression is that the reluctance of many authors to accept anything that appears to be contrary to methodological or ontological individualism comes from not distinguishing the two separate notions of deontic status and psychological state. We think it is reasonable to deny the existence of “collective minds,” “collective emotions,” and so on, as long as psychological states are always states of some individual’s brain,<sup>13</sup> though, obviously, individual brains can be as densely interconnected as you like to others through social and cultural mechanisms. But this reluctance has been illegitimately extended to the denial of “collective beliefs,” “collective intentions,” “collective plans,” or “collective reasoning,” by considering that beliefs, intentions, and so on, are necessarily identical or reducible to psychological states or processes. Our inferentialist approach allows to see, instead, that there are two different things that contribute to making of something a (rational) belief: some psychological states, on one hand, and the subjection to some inferential norms of the deontic statuses those states institute, on the other hand. The important point now is that *the subject whom the deontic status belongs to is not necessarily the same one as the subject whom the psychological states instituting that status belongs to*; this will depend on the inferential practice in which the belief, intention, and so on, is framed. This means that it is not necessary that a collective subject has anything like “collective states of mind” to have, for example, beliefs *in the sense of claims the subject is responsible for*. Actually, it is not even necessary that one collective entity’s status of having a certain “belief” (e.g., an economic research institute publishing some forecasts) is *immediately* caused by some individuals having certain mental states in that precise moment (in the same example, imagine the forecast is automatically produced by a computer program from numbers that are automatically downloaded from other databases). Of course, some individual psychological states must exist *at some or other point* for the collective entity to have a particular practical or doxastic commitment. This is so because, at least as far as we know, the sort of flexible and sophisticated behavior required to be counted as taking part in a normative practice can only be produced by the psychological processes that take place in human brains. But the inferential rules that institute those collective normative statuses can “automatize” the process in such a way that there can be a considerable spatial, temporal, and institutional distance between the actions of the individuals and the ensuing actions of the collectives.

Some could be tempted to generalize the possibility we have just exemplified, by asserting that individuals have something like “original” normative statuses, whereas collectives can only have “derivative” ones. Perhaps this is true, though we think that this is not a metaphysical, but an empirical (though extremely difficult to answer) question. In the first place, even if it is true that

a psychological state is always the state of an individual's mind, this is far from being a decisive argument in favor of "ontological individualism," at least in our comprehension of how our cognitive capacities work; for, exactly in the same way as human infants are obviously not "nutritionally self-sufficient," they are also not *cognitively* self-sufficient: the construction of an individual mind, capable of having rational, discursive commitments, is a process in which the interaction with a stable social environment of cognitive tools, norms, and practices is absolutely essential, and this environment is constructed only through the continuous interaction, negotiation, competition, and cooperation between hundreds, thousands, and, in modern societies, millions of individuals, and this is far from being wholly replicable in the mind of one single human being. Individuality is more a question of where the locus of *responsibility attribution* lies (and hence a normative trait) than of physical or cognitive *self-sufficiency* (which, in the case of concrete organisms, is always more or less illusory). Furthermore, it is far from clear whether the causal and inferential processes by which an individual enters into some deontic commitments can *always* take place without assuming some previous collective commitments, that is, whether these processes can be described as being *exclusively* composed of normative statuses and psychological states of individuals. Certainly, it seems that Brandom's discursive practices can be characterized without explicit appeal to collective commitments. However, in Brandom's proposal, individual commitments are instituted through complex collective interactions: they presuppose the existence of *social* practices whose members are subject to inferential norms that affect the whole community. After all, one's learning an inferential norm (e.g., the meaning of a word) seems to consist in acknowledging that a particular normative practice (the rules telling what uses of the word are correct) is *collectively* accepted within a particular community, and as each individual can to a certain extent "misrepresent" the meaning of the word, perhaps (and perhaps not) the best way of describing this learning process is by saying that there is a kind of "consensual meaning" to which no individual member of the community perfectly adapts, but which has the status of a "implicit" communal commitment. So, the question of what comes ontogenetically first, individual or collective normative statuses, is probably a chicken-and-egg problem whose solution we do not dare to conjecture, and that is probably irrelevant to understand the *actual* working of normative practices (which function thanks to the actual existence of both individual and collective deontic statuses), as knowledge about the origin of life is not necessary to understand to a high extent most biological phenomena.

So, our argument can be summed up as follows. According to inferentialism, attributions of intentional states are understood as attributions of

normative statuses—of commitments and entitlements. Therefore, when we attribute a belief or an intention to some entity (to an agent or system) we do not need to attribute a mind or a psychological process to that same entity. In this way, collective intentionality will be possible in practices in which commitments and entitlements can be attributed to, and undertaken by, collectives or groups. For this to happen, certain performances have to be attributable to a collective as a whole: these performances will be interpreted as performances for which the collective is held responsible—and as a result of some of them, the collective will count as undertaking a commitment. Collectives or groups, consequently, will be assessed with regard to these performances by the other participants in the practice, and may be sanctioned in accordance (for instance, through the removal of previous entitlements, or with some form of direct punishment, like a financial fine).

The conclusion of this section conveys the fundamental message of our article: commitments and entitlements, both individual and collective, are *neither psychic nor physical states*, but normative statuses, instituted by their being attributed or acknowledged by some agents according to some inferential norms endorsed by those agents (who are the relevant agents and what are the relevant norms will depend on the structure of each particular practice). To this extent, collective commitments and entitlements pose no harder ontological problems than individual ones: they are simply states of the dynamic *score* of a game; they are, so to say, “mere software,” and it is indifferent *for their function as scores* what is the biological, psychological, or artificial nature of the “hardware” on which they are “run” (though, in fact, only a mature and culturally sophisticated group of individuals can provide the right “hardware” for the emergence of *both individual and collective commitments*).

## 5. Conclusion: Collective Agents as Virtual Prostheses

In the previous sections, we have introduced an image that is useful for illuminating our view of collective agency: the image of agency as a kind of virtual machine, that is, a kind of “software” that can be “implemented” on very different kinds of “hardware.” The software constituting this machine would not consist of series of numbers, as in the case of computer software, but of systems of norms and normative statuses: to be a rational agent is to be able of being in normative statuses, that is, being capable of behaving in a way that can be interpreted as following *inferential norms*. Sane human beings are *natural* agents, that is, their biological hardware is naturally

prepared to behave in such a way, but *artificial* agents can also exist, that is, systems that have been created by people so that they are able to behave in response to reasons. The “hardware” of these artificial agents would usually consist of groups of human individuals plus every other material element necessary for the proper functioning of the artificial agent, but their “software” would also consist, as in the case of individual agents, of systems of normative statuses connected by inferential norms. An important clarification is that, though in cognitive science it is usual to appeal to the metaphor of “brain as hardware versus mind as software,” what we are suggesting is that the *mental or psychological states* of individual human beings, as well as their *causal* psychological connections, had to be taken rather as part of the individual agent’s “hardware,”<sup>18</sup> whereas the “software” would be the *normative import* of those states. The difference is clearly seen when we consider the difference between the *mental state* of being conscious of believing something (e.g., being in the conscious state describable as “I have to pay this bill tomorrow”), and being *committed* to that something: there are cases in which you may have the obligation of paying the bill even if you have forgotten it, and even if you have never been conscious of having acquired that obligation. As we have seen before, your normative status does not consist only in the commitments you actually acknowledge but also in those commitments that *follow* from your previous actions plus the norms you are subject to. Our claim is, hence, that *the normative statuses of collective agents do not need to be underlain by something like “mental states” to count as normative statuses*. A paradigmatic example would be a couple of big firms (e.g., an electric company and a car manufacturer), one of which issues in a purely automatic way an invoice to the second one, which in its turn pays also automatically through an electronic bank transfer. Perhaps no individual employee of any of both companies (nor of the banks involved) happens to notice this transaction, but this is no reason to deny that there has been some change in the normative statuses of the firms, as well as some processes that can be interpreted as their *claims* and their *actions*.

From this point of view, a collective agent could be compared with a kind of artifact, a “virtual prosthesis,” an artificial entity governed by the dynamics of inferential norms, and that people create to be capable of doing things that are more difficult, and often impossible, to be performed just by

---

<sup>18</sup>Note that we are talking of mental and psychological states in broad terms, without assuming that they consist of representations. In this article, we are not committed to any specific view of the “hardware” underlying normative practices.

individual human beings working in isolation.<sup>19</sup> For a collective agent to exist, thus, it is only necessary that people establish some inferential rules that allow to determine the agent's normative statuses and dynamics, that is, what facts (most of them performable by individual agents, but not necessarily all, as we have seen in the case of automatic invoicing) will count as the acquisition of commitments or entitlements by part of the collective agent, and what changes in its normative statuses or those of other agents would the actions of the collective agent entail. As all other kinds of *artifacts*, the range of actions that a specific collective agent will be able to perform is limited, in this case not merely because of the physical properties of the systems the agent employs (as a violin cannot be used to move a car), but also because of the restricted scope of the norms that make of this entity an agent (e.g., a firm cannot be married). But this limitation only reflects the fact that collective agents are artifacts with a limited number of purposes, and does not entail that they are not "real" agents, for, to count as agents, it is only necessary that they are able to behave in a way that is consistent with the normative statuses they have, not with normative statuses we are not entitled to attribute to them.

This view of collective agents as artificial normative entities allows an easy comparison of our approach with other theories of collective agency. In the first place, our approach is agnostic on the question whether collective minds and mental states exist or not,<sup>20</sup> we only claim that they are *not necessary* for collective agency. What is necessary, instead, are collective *normative statuses*, that is, commitments or entitlements that pertain to no individual human being as such, but to a group or institution that can be logically and normatively separated from its individual members. Of course, if collective mental states exist, they can in a natural way play the role of (part of the) "hardware" in those collective agents whose normative statuses are their "software," as individual mental states do in the case of individual agents, but what our approach shows is that the question of individual versus collective *agency* is orthogonal to the debate about the ontological status of collective *psychological states*. So, one can be an ontological or methodological individualist, in assuming that only individuals have psychological states, but simultaneously accept without any kind of metaphysical worry the existence

---

<sup>19</sup>This view, of course, is cognate to the idea of scientific models as "inferential protheses" defended in Donato Rodríguez and Zamora Bonilla (2009).

<sup>20</sup>For a defense of the possibility of collective mental states, see, for example, Brooks (1986), Velleman (1997), or Wilson (2005).

of collective agents and collective agency, for these can be taken as mere systems of attribution of normative statuses backed by some system of behavior that can be interpreted as *fulfilling* what is commanded or permitted by those statuses. The question whether collective agents are “real” or “fictitious” is pointless, because it is not really a *philosophical* question for the metaphysician, but a *practical* question for real people, who need to know whether in certain circumstances *they* (i.e., the people) are obliged or entitled to accept claims of the form,

this collective agent (e.g., a firm, a church, a country . . .) has such and such commitments or entitlements, can is entitled to do such and such things in order to comply with them, and this has such and such implications for my own commitments and entitlements.

Stated in other way, the only relevant question is not whether the attribution of collective commitments is “metaphysically true,” but whether it is *normatively appropriate* for a specific agent under specific circumstances.<sup>21</sup>

Individualist philosophers have tended to escape from the necessity of accepting collective agents by reducing them to some particular kind of psychological “mode” in the individual minds. Typical examples of this approach are the views of Bratman, Gilbert, Searle, and Tuomela.<sup>22</sup> Although these authors disagree about many things, what is relevant from the point of view of the question we are discussing is the fact that for them there is not any collective mind, but joint intentions that ultimately consist in a specific way of *individual* minds being in themselves (in Tuomela’s terms, thinking in the “we-mode”) and being *interrelated* among themselves. Of all these four authors, Gilbert is probably the one who has most perspicuously insisted on the importance of the *normative* nature of the individual attitudes in constituting the collective intentions. From the point of view of game theory, Robert

---

<sup>21</sup>Collective rational agents—that is, collectives or groups that count as taking part in a discursive practice—must be able to follow the norms of the practice with some degree of success. In particular, their collective behavior must show, as we indicated for individual agents, some sensitivity to the evaluations and sanctions of the other participants; moreover, there must be some minimal coherence in the commitments acknowledged by the collective (and in their relation to the collective’s actions). Otherwise, the collective will not be regarded as a unified subject capable of coherent rational agency.

<sup>22</sup>See, for example, Bratman (1999), Gilbert (2000), Searle (1995), and Tuomela (1995).

Sugden<sup>23</sup> has similarly proposed the idea that individuals can shift from “individual utilities” to “team utilities,” though it is obscure how and why they can do the “shift”; strategic thinking with team preferences would lead to team reasoning, but it is also reasoning performed by individual minds, however interdependent they are. Our approach is much more similar to that of Schmid (2009), who also argues that collective intentions are essentially not psychological, but normative facts, and to that of List and Pettit (2011), who claim that for a group to be an agent it suffices that it satisfies three “conditions for agency”: having representational states, motivational states, and the capability of processing those states so as to intervene in the world; rationality would also demand that the agent attempts to keep some degree of coherence between its multiple cognitive states, and between these and its actions. Our approach can be understood as an attempt to reach a similar conclusion about the reality of group agency starting from an inferentialist (i.e., nonrepresentational) view of rationality. As List and Pettit explain, the most important problem for attributing agency to groups is the question of how to “hold them responsible.” This entails the group’s capacity of engaging into a process of argumentation with other agents (something for which we think an inferentialist theory is better prepared), and, as the same authors argued in other works, it can be the case that there is no way of aggregating the group’s members individual attitudes in a consistent way as to constitute a coherent group’s attitude (or set of doxastic commitments), so that the latter is not mechanically “reducible” to the former, that is, the group must have a set of (as we would say) inferential constraints proper of it and different of those of each individual member.<sup>24</sup> To the extent that the group achieves this capacity, it would have, in Pettit’s terms, “a mind of its own,”<sup>25</sup> though, as we have seen, we prefer to reserve the word “mind” for the psychological states, rather than for that space of inferential commitments that can be attributed both to individual and collective agents.

In conclusion, we think that inferentialism, as a theory of intentionality, offers the most illuminating approach to the study of collective intentional states. This is so because, as we have seen, inferentialism avoids identifying intentionality with the possession of some specific psychological-neurological states by part of the agent. To count as a rational agent, it is enough that a

---

<sup>23</sup>See, for example, Sugden (2000).

<sup>24</sup>See Zamora Bonilla (2007) for an explanation of how some of constraints of individual judgment aggregation can give rise to rational collective sets of judgments.

<sup>25</sup>See Pettit (2003).

system is able to play a suitable role in discursive normative practices—in particular, that such a system is able to undertake certain normative statuses. Thus, an inferentialist approach to intentionality allows us to explain the attribution of intentional states to collective entities without needing to presuppose the existence of collective brain states or psychological processes (even if it might turn out to be necessary that the individuals that compose the collective possess some psychological states). Perhaps there are other approaches to intentionality and rational agency that yield similar results, but we think that inferentialism manages to do so in a remarkably straightforward and unproblematic way.

### Acknowledgments

For comments and discussion, we would like to thank Daniel Whiting and audiences at Collective Intentionality VIII (Manchester), the III Madrid Workshop on New Trends on the Philosophy of the Social Science, and the 2013 meeting of the Spanish Society for Analytic Philosophy.

### Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

### Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was supported by the Spanish Government research project FFI2011-23267 (“Inferentialism as social epistemology”). Javier González de Prado’s work was also supported by the FPI-UNED doctoral scholarship.

### References

- Bouvier, A. 2011. “Individualism, Collective Agency and the ‘Micro-Macro Relation.’” In *The SAGE Handbook of Philosophy of Social Sciences*, edited by I. Jarvie and J. Zamora-Bonilla, 199-216. London: Sage.
- Brandom, R. 1994. *Making it Explicit*. Cambridge: Harvard University Press.
- Brandom, R. 2008. *Between Saying and Doing: Towards an Analytic Pragmatism*. Oxford: Oxford University Press.
- Bratman, M. E. 1999. *Faces of Intention*. Cambridge: Cambridge University Press.
- Brooks, D. 1986. “Group Minds.” *Australasian Journal of Philosophy* 64:456-70.
- Donato Rodríguez, X., and J. Zamora Bonilla. 2009. “Credibility, Idealisation, and Model Building: An Inferential Approach.” *Erkenntnis* 70:101-18.
- Gilbert, M. 2000. *Sociality and Responsibility: New Essays in Plural Subject Theory*. Lanham: Rowman & Littlefield.

- Grönert, P. 2005. "Brandom's Solution to the Objectivity Problem." *Pragmatics & Cognition* 13:161-75.
- Hattiangadi, A. 2007. *Oughts and Thoughts*. Oxford: Oxford University Press.
- Jarvie, I., and J. Zamora-Bonilla. 2011. *The SAGE Handbook of Philosophy of Social Sciences*. London: Sage.
- Kaufmann, L. 2011. "Social Minds." In *The SAGE Handbook of Philosophy of Social Sciences*, edited by I. Jarvie and J. Zamora-Bonilla, 153-81. London: Sage.
- Kuklla, R., and M. Lance. 2009. "Yo!" and "Lo!": *The Pragmatic Topography of the Space of Reasons*. Cambridge: Harvard University Press.
- Laurier, D. 2005. "Pragmatics, Pittsburgh Style." *Pragmatics & Cognition* 13:141-60.
- List, C., and P. Pettit. 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- MacFarlane, J. 2010. "Pragmatism and Inferentialism." In *Reading Brandom*, edited by B. Weiss and J. Wanderer, 81-95. New York: Routledge.
- Pettit, P. 2003. "Groups with Minds of Their Own." In *Socializing Metaphysics*, edited by F. Schmitt, 167-95. New York: Rowman & Littlefield.
- Rödl, S. 2010. "Normativity of Mind versus Philosophy as Explanation." In *Reading Brandom*, edited by B. Weiss and J. Wanderer, 63-80. New York: Routledge.
- Rosen, G. 1997. "How Makes the Rules Around Here?" *Philosophy and Phenomenological Research* 57:163-72.
- Schmid, H.-B. 2009. *Plural Action*. Dordrecht: Springer.
- Searle, J. R. 1995. *The Construction of Social Reality*. New York: Penguin Books.
- Shapiro, L. 2004. "Brandom on the Normativity of Meaning." *Philosophy and Phenomenological Research* 68:141-60.
- Sugden, R. 2000. "Team Preferences." *Economics & Philosophy* 16:175-204.
- Tuomela, R. 1995. *The Importance of Us: A Philosophical Study of Basic Social Notions*. Stanford: Stanford University Press.
- Velleman, D. 1997. "How to Share an Intention." *Philosophy and Phenomenological Research* 57:29-50.
- Wanderer, J. 2008. *Robert Brandom*. Montreal: McGill-Queen's University Press.
- Wilson, R. 2005. "Collective Memory, Group Minds, and the Extended Mind Thesis." *Cognitive Processing* 6:227-36.
- Zamora Bonilla, J. P. 2007. "Optimal Judgment Aggregation." *Philosophy of Science* 74:813-24.

## Author Biographies

**Javier González de Prado Salas** is a graduate researcher at UNED, Spain, and the University of Southampton, UK. He works on pragmatist approaches to intentionality, in particular on inferentialist accounts of representation. He is also interested in the relations between inferentialism and contextualist perspectives on thought and language, and, especially, in how this sort of views may provide a suitable account of metaphor and other aesthetic phenomena.

**Jesús Zamora Bonilla** is a full professor of philosophy of science at UNED, Spain. He has worked on philosophy of social sciences, scientific realism, and the economic approach to epistemology. Some recent publications are “Why are good theories good?” (*Synthese*, 2013), “Conversation, Realism, and Inference” (in *Economics for Real*, Lehtinen et al., eds., Cambridge University Press, 2012). He is also coeditor of *The SAGE Handbook of Philosophy of Social Sciences* (SAGE, 2011, with Ian Jarvie).