

# **The Normativity of Group Agents: Expanding the Organisational Account**

Jimmy Lewis-Martin

jimmy@uow.edu.au

## **1 Introduction**

Recently, there has been much discussion of the norms that drive collective systems such as businesses, political parties, universities, and charity organisations (List & Pettit, 2011; Pauer-Studer, 2014; Tollefsen, 2002; Tuomela, 2013). Such groups dominate our social and political landscapes, their activities dictating our legal structures, the availability of education and healthcare, and our collective leap into climate crisis. The behaviours of many of these groups are best explained by ascribing agency to them since it is explanatorily indispensable for understanding their behaviour (List, 2021, p. 4). So, having an account of group agency is central to coming to grips with the systems that influence our own lives.

An interesting difficulty encountered with describing groups as agents is that there seem to be many instances in which the activities of a group go directly against the interests of its members. Choosing more environmentally harmful methods of resource extraction in order to maximise profit, for instance, is obviously detrimental to the people involved but is beneficial to the group. Sales targets in many workplaces are likewise harmful to the psychological wellbeing of employees but are excellent for pushing workers to their limits in service of the group agent. In eusocial non-human group agents, as Mog Stapleton and Tom Froese (2015) point out, the individual members will often act to benefit the group to their own direct detriment, as when a bee stings a person at the cost of its own life (p. 226). If we wish to understand group behaviour, especially where these groups act contrary to our expectations, where they depart from the desires of their members, and where they continue to pursue harmful actions despite legal or social sanctions, then it is crucial that we understand both the

norms of these group agents and how these norms arise. As I will argue below, it is the normativity of group agents that determines which actions they choose over others. Furthermore, where a group's norms come from makes a difference to how we deal with unjust collective actions. If they are dictated by one person, then we can replace that one person. If they are instead determined by the group's organisational structure, then structural changes may be more appropriate.

To address the dual questions of what a group's norms are and where they come from, I will argue for applying and expanding the *organisational account* of normativity to group agents. The organisational account finds its roots in the works on biological systems of Georges Canguilhem (1991) and Hans Jonas (1966) and has gained a good deal of traction in the enactive theory of agency (Barandiaran et al., 2009; Di Paolo et al., 2017, pp. 120-4). For the enactivists, agency involves at least a *system* that *acts* on its environment according to certain *goals* or *norms* (Barandiaran et al., 2009, p. 369). From this basic outline of agency, Xabier Barandiaran, Ezequiel Di Paolo, and Marieke Rohde (2009) identify three central features of agency: individuality, interactional asymmetry, and normativity (p. 369). Importantly, in each case a core concern involves understanding the agent on its own terms. That is to say, it is not enough that something looks like a coherent system to us or that it seems to be pursuing goals from our point of view. To be a satisfactory account of normativity, then, our theory must both identify the norms of a given agent and it must explain how those norms arise for or from the agent itself. Understanding how agential norms occur is therefore both practically important, as I argued above, and methodologically important, as the enactive theorists claim.

I will show that the organisational account of normativity as it is expressed by Canguilhem (1991) and built on in the enactive literature explains the complex nature of normativity in group agents. Roughly, advocates of the organisational account maintain that the norms of an agent are determined by that agent's organisational structure insofar as that

structure dictates both what the agent is capable of and what is necessary for its persistence in its present environment. It is the *organisational* account since it holds that the generation of norms is relative to the agent's structure, rather than in the ideas of particular individuals or written statements or anything else.

I will start by explaining the minimal version of the organisational account, which is the account as it is found in Canguilhem and the enactive theory of agency. It is the minimal version of the organisational account because it is concerned only with the norms that every agent must have to be a successful agent. I then show that this minimal organisational account can form the basis for understanding normativity in collective systems. To develop a complete view of the norms of group agents, it will be crucial to build on the minimal organisational account to address the specific, non-minimal aspects of group agency. These are additional norms that the group might adopt as a result of being made up of other autonomous agents as well as the enforcement of the group agent's own norms over those of its members. Having shown that the organisational account is feasible, I compare it to some prominent views in the social ontology literature, namely Deborah Tollefsen's (2002) *interpretationism* and Herlinde Pauer-Studer's (2014) *constitutivism*. In doing so, I aim to show that the organisational account is both viable and is preferable to the other available accounts in the literature.

## **2 The Organisational Account**

To argue for the organisational account, I will trace roughly its historical progression from Canguilhem to the contemporary enactivists. This is because Canguilhem introduces most clearly the central ideas that guide the account, which are then built on by the enactive theorists. The historical progression of the account and a concise explanation of it therefore coincide. I will begin by briefly sketching the argument to follow so that the relevance of each point below is clear.

Canguilhem (1991) introduces the idea that agents establish their own normativity in response to the demands of their environment (pp. 126-7; p. 143). Canguilhem, however, does not give a precise explanation of *how* this process occurs (Barandiaran & Egbert, 2013, p. 9). Barandiaran and Matthew Egbert (2013) point to Jonas for the solution (p. 9). Jonas argues that the identity of an organism is given by its organisation, not by its constituent material components (1966, p. 78). It is this functional organisation that determines the norms that can arise as a result of the agent's interactions with its environment. What is required in a given environment for the system to be viable is determined by the structure of the system itself. Finally, the concept of *adaptivity* is introduced to explain how particular actions can count as good or bad according to the system's own normativity (Di Paolo et al., 2017, p. 122). An action is adaptive if it takes the system further away from failure. Importantly, this can include the agent affecting the environment around it so that the environment is more conducive to the agent's needs, as with agriculture. Minimally, then, an agent's normativity is given by its conditions of viability, which are themselves determined by the agent's present organisational structure interacting with its present environment. That animals require food, for instance, is a requirement of their particular organisational structure. How they get it depends on the capacities afforded to them by that structure and the availability of food sources in their environment. A particular action is normatively good or bad depending on whether it increases or decreases the agent's chances of achieving its normative ends (getting food) or maintaining its normative conditions (surviving).

The organisational account finds its first expression in Canguilhem (1991). The account of normativity he gives is one that is essentially generated by the system's interactions with its environment. 'Taken separately, the living being and his environment are not normal: it is their relationship that makes them such' (Canguilhem, 1991, p. 143). To make his point, Canguilhem (1991) points out that in industrial areas in England and Germany, black butterflies replaced

grey butterflies of the same species (pp. 142-3). The black butterflies, he claims, were typically far more vigorous than their grey counterparts. In their usual habitat, the black butterflies stand out too much against the bark of trees and are easily predated upon. In industrial areas where there are fewer birds to threaten them, however, the black butterflies were able to thrive (Canguilhem, 1991, p. 143). Different environments, in other words, will favour different solutions. In forested areas, the less vigorous but more well-camouflaged butterflies outcompete their darker relatives. When there are less predators, being more vigorous is advantageous. ‘A living being is normal in any given environment insofar as it is the morphological and functional solution found by life as a response to the demands of the environment’ (Canguilhem, 1991, p. 144). Normativity arises out of the interaction between a particular system and its environment. Put simply, this is necessarily the case since action always occurs within an environment. What an agent needs and how an agent can get what it needs are always present outside of itself, since otherwise the agent would not need to act to attain those things at all. Likewise, how the agent must behave or respond to the demands of its environment will change depending on environmental context.

Canguilhem’s insight has been productively employed by enactivists, who build on his view of normativity by connecting it more explicitly with the organisational structure of agential systems. Taking Canguilhem’s account as it stands, we have an account of what agential norms are, but not a full view of how they occur. Canguilhem rightly points out that norms arise via the agent-environment interaction, but he does not explain how the agent influences this interaction. Furthermore, likely because of his focus on the health and pathology of the system as a whole, Canguilhem’s discussion is not overly concerned with normatively good or bad *behaviours*. The enactivists resolve these issues while maintaining the core of Canguilhem’s arguments by drawing on the work of Jonas (1966) and by introducing the concepts of *adaptivity* and *sense-making*.

According to Barandiaran and Egbert (2013), detailing the occurrence of norms was an important contribution of Jonas (1966). Jonas argues that living systems cannot possibly be identified strictly as their present material composition, since organisms are in a constant state of self-construction and renewal. Hence, the organism must be identifiable instead as a continual process – a functional system that constantly takes new material into itself in order to persist as it uses up or loses old material. As a result, the organism necessarily depends on its own continual activity for its existence. The normativity of the living system is a necessary feature, since to be a living system is to be an active system that must pursue its own self-maintenance.

Jonas (1966) adds that the capacities of an agent coincide with its necessary behaviour (p. 105). Animals gain motility over plants and in that sense gain a new capacity. With this freedom of motility, however, comes the requirement for motility. Without the vital support of the soil that plants are attached to, animals must make use of their additional freedom of movement in order to meet their needs. The organisational structure of an agent determines both what an agent is capable of doing and what it must do to maintain itself.

To address the normative thrust of particular actions or behaviours of an agent, Di Paolo et al. (2017) argue that agents are *adaptive* systems (Di Paolo et al., 2017, p. 122). Adaptivity involves more than merely avoiding death, but rather requires the capacity to differentiate between beneficial and harmful outcomes relative to the system's viability and acting to improve that viability as well as avoiding or addressing threats (Di Paolo et al., 2017, p. 122). Adaptively successful systems will be far more capable of absorbing the cost of mistakes than systems that exists in a closer-to-failure state. Nevertheless, given the normative pull of adaptivity, we can still consider maladaptive behaviours in otherwise successful systems a normative failure. Without the concept of adaptivity, however, the system only fails when it

ceases to be viable. Given that viability guides adaptivity, however, we can consider actions good or bad as they take the agent further away or closer to failure.

Adaptivity is required for *sense-making* (Di Paolo et al., 2017, p. 123). As Di Paolo and colleagues (2017) put it, ‘by “sense-making”, then, we refer to the notion that objects or events become meaningful for an agent if they are involved in the normatively guided regulation of the agent’s activity’ (pp. 123). Sense-making, then, is essentially the process of experiencing things as meaningful. This process is guided by the system’s own normativity. Small animals like birds and cockroaches will typically flee if a human approaches them. Even if the human has no intent to harm, the fleeing behaviour is an adaptive one because the human always could, on the basis of its size alone, cause harm to the smaller animals. Likewise, bears are not always aggressive to humans, but a bear in the wild is meaningfully perceived as threatening. No reasoning is involved in the process of perceiving the threat – it just is in virtue of its evident potential to cause harm.

It is the explicit embedding of normativity in the *structure* of the system as it interacts with its environment that primes the organisational account for expansion to group agents. Like organisms, the constituent material parts of groups are constantly being renewed, leaving the series of functional processes that make up the system’s organisational structure as the only reference identity. Likewise, groups are necessarily actively oriented toward their own viability, which is determined by its functional structure in the environment it is in. Particular actions will be good or bad for the group as they take it further from or closer to a failure state. For instance, businesses pay attention to market trends, which form a core part of their environment, and spend a good deal of money on research and development not just so they can break even, but in order to further maximise their profit margins. This, too, is the point of sourcing cheaper materials and labour. It may still be profitable to use more expensive materials, but it is adaptively beneficial for the business to pay less and achieve more significant

profits. Doing so is only possible so long as cheaper materials and labour are available to the business. Thus, they may also lobby the government to ensure that no laws are introduced to prevent their outsourcing – an instance of an agent influencing its environment so that it can remain adaptively viable. Evidently, then, group agents are easily described as organisational structures actively oriented toward maintaining their own viability given the demands of their present environment. The organisational account readily explains the norms of group agents.

It would be contentious to claim that group agents themselves have sense-making capacities, though we can think of their activity as being guided in a similar way. The members of the group can perceive the environment as meaningful for the group in their capacity as members. Sometimes, an individual's entire role may be centred around performing 'sense-making' tasks for the group, as with analysts checking over feedback after product testing. Likewise, algorithms are designed to differentiate between 'good' and 'bad' according to particular, built-in standards. Whether or not this means the group itself has sense-making capabilities is deserving of its own investigation. Nevertheless, they do achieve something functionally analogous.

The minimal version of the organisational account that I have described here is what allows further norms to be formed at all, since if the agent does not achieve its minimal normative requirements, the agent will cease functioning. With this in mind, I will build on the organisational account of normativity to better address the development and pursuit of norms in groups.

### **3 Normativity in Group Agents**

So far, I have outlined the minimal version of the organisational account. It is necessary that a system self-maintains if it is to have norms at all. However, as we know from experience, there can be layers of norms beyond this minimal level, created or adopted from different



sources. For us, there are social norms, legal norms, professional norms, norms that arise from personal goals that we have set, which have themselves been set for all sorts of normative reasons, and countless others. Group agents that have been formed by humans are similarly capable of various kinds of norms, so long as its minimal normative requirements are met. Here, I will expand on the account to better address the complexities of group agency, first by discussing the interactions between a group's minimal norms and its other norms and then by looking at the enforcement of a group's own normativity, especially where it conflicts with the norms of the group's own members.

### 3.1 Extra-minimal norms

Group agents are not strictly minimal agents because their normativity is not dictated *entirely* by their self-maintenance. Instead, because groups are often made up of people with their own ideas and norms unrelated to the group's own minimal normativity, they may become the loci of various other norms that are either unrelated or even detrimental to the group's viability. A private fast food chain might close its stores on Sundays because of the religious views of its founder(s) and executive members, for example. A clothing company might also be willing to pay more for its materials if it is ethically sourced at the cost of a slight increase in profits. These are the kinds of norms that I mean by 'extra-minimal'.

In these instances, their drive to self-maintain acts as a motivating and constraining force on the group's members and its potential norms. It motivates because the system is structured as a self-maintaining system, meaning the members of the group will be members precisely because they hold a position that involves achieving the group's end, whether directly or indirectly. By virtue of one's position in the group, they are already primed to act in the group's interest, i.e., to do their job. It constrains because it limits the possible norms that can be adopted by the group, since it cannot adopt norms that harm its adaptive capabilities too

drastically, though some reduction in viability can potentially be tolerated depending on the robustness of the group. So, the clothing company that pays more for its materials will make less profit than it otherwise could, but if the company is successful enough it can afford to commit itself to this moral norm even though it may be more viable without it.

This capacity for adopting norms that are not strictly dependent on the system's viability is an inherent aspect of relying on other autonomous agents for one's own material constitution, especially when those autonomous agents are human beings who themselves are the loci of various extra-minimal norms. That said, there will always be pressure to purchase cheaper, less ethically harvested or produced materials or to outsource to more heavily exploited workers because of the chance at increased profits, which would be adaptively beneficial for the agent. Members in management positions may even be situated due to their position to see this adaptive benefit as meaningfully important for the group. Hence, we should always expect worsening moral conditions inside a group agent unless its environment is sufficiently changed to make this unlikely.

The organisational account applied to group agents, then, holds that group norms are generated in the first place by their conditions of viability. However, within the bounds of the group's capacity to remain viable, there is space for additional norms. Groups can be viable without resorting to sweatshop labour – though they may be less robust, i.e., less capable of persisting through poor business decisions or detrimental environmental conditions like economic recessions.

The organisational account leaves room for additional kinds of normativity in group agents. These other norms will always be constrained by the group's capacity for self-maintenance. It is this constraint that ultimately explains why groups tend toward acting in their own interest to the detriment of people without this universally being the case.

### 3.2 Group norm enforcement

One particular problem that confronts us when considering the norms of group agents is unjust collective behaviours enacted by people who are at least personally opposed to those behaviours or, in more extreme cases, are personally harmed by them. A complete theory of group agential normativity therefore requires some account of how group norms are enforced. Fortunately for groups, there are already a multitude of factors working in their favour. First of all, humans are fundamentally social animals, so we tend toward creating and acting as parts of groups. Beyond that, there are economic, social, political, and psychological pressures on individual people that will leave them more or less willing to participate in the activities of group agents, no matter what they are. For example, the threat of poverty and the allure of wealth are both important economic factors in coaxing participation in collective actions out of people. Expectations from friends and family, the desire to be perceived in certain ways by other people, and our own sense of achievement and self-worth are all other, perfectly plausible avenues through which pressure is applied on individuals to participate in collective actions.

Beyond these, norms are enforced partially through the group's structure via the positions that people hold. Simply put, if a person takes up a role in a group, they have tacitly agreed to perform the relevant functions of that role. The agreement in question may simply express that they are willing to do their job for the compensation offered, perhaps to avoid poverty which they might perceive as being worse than whatever they are asked to do on behalf of the group. The participation in these roles by people who would be opposed to the group's activities on moral grounds is made easier because of the division of labour and the possibility that group members can be blind to the broader goals that they are working toward.

Finally, ideology may play an important and pervasive role in ensuring compliance. The notion that everyone has a duty to contribute to society, for example, is especially prevalent. It is common to hear the question 'what do you do [for work]?' when meeting

someone new, and it feels embarrassing to say ‘nothing’ in response. Referring to working as ‘contributing to society’ similarly embeds the notion that work is an inherent good, along with the derisive language used to describe people who receive welfare payments. All of this, among countless other examples, is enforced via the actions of political parties, reflected in the language they use and the policies they endorse.<sup>1</sup>

Each of these factors – the economic, social, political, and ideological – form an essential component of the environment for individual people. At the present moment, these factors likewise coincide such that submission to the demands of groups is the default.

These points all speak to norm enforcement in group agents whose constitutive activity is in some way undesirable, unjust, or unfulfilling for the individuals pursuing those collective ends. Of course, not all group agents are so horrid, as is evidenced by the general success of charitable organisations, local sporting clubs, or other groups oriented around other fulfilling or morally upstanding activities. There is less mystery surrounding member participation in these kinds of group agents. If the members perceive the group and their actions in it as good, even when they have in view the full normativity of the group, then it is much simpler to understand, psychologically, why people would participate. Furthermore, groups are very rarely wholly good or bad, and so both personal desire and external pressure is often involved. Working in academia is desirable to some people because they enjoy the research or teaching. It is perfectly coherent to accept that these same people think the academy itself is fundamentally flawed, that the university has become commodified, and that abusive people abound and are readily protected by their institutions and colleagues. In these cases, the people involved are both willingly participating in particular processes within their respective group

---

<sup>1</sup> In Australia, it is mandatory to apply for a certain number of jobs and to go to an employment agency as often as they ask you to in order to receive welfare payments. The goal of the agency is to find you employment, no matter what, and it is common to be pressured into applying and interviewing for jobs that are entirely unrelated to your education, expertise, or preferred field of work.

agents and are either helping enact harmful processes or at least failing to undermine those processes as a result of the pressures discussed above.

The aspects of group norm enforcement discussed here help to explain how and why people participate in the actions of groups that they might consider wholly or partially bad. Even in cases where we might think the group itself has noble ends it is possible that we participate in patterns of behaviour that we ultimately would prefer did not exist.

#### **4 Comparative Discussion**

Now I will compare the organisational account to some prominent existing views of collective normativity, namely those of Tollefsen and Pauer-Studer. I am focusing on these two accounts in particular because they both aim to directly address the question of what a group agent's norms are. Though there are many other discussions related to the normativity of group agents, these are not always concerned explicitly with determining what the norms of group agents are. Raimo Tuomela (2013), for example, is often concerned more with the roles of a group's members in accepting and achieving a group's goals rather than strictly on the group's normativity itself. In other cases, we are given minimal constraints on a group's norms, but not told how to determine what they are (Rovane, 2019, pp. 4872-3). Although these discussions may be partially related to the one here, I have chosen to focus instead on arguments that are directly concerned with the same issues that I am.

Tollefsen (2002) argues that *interpretationism* as developed by Donald Davidson and Daniel Dennett allows us to interpret the actions of groups, which for the interpretationist means that those groups really are intentional agents (p. 397). Interpretationism involves taking the intentional stance toward the target of explanation and providing reasons for their behaviour (Tollefsen, 2002, p. 398). Importantly, these reasons must justify the behaviour from the agent's own perspective (Tollefsen, 2002, p. 398). Tollefsen cites William Taschek to help

make her point: it ought to make ‘sense *to us* that it made sense *to the agent* that she did what she did’ (Tascheck in Tollefsen, 2002, p. 398). In so doing, we must assume that the agent in question shares our norms of rationality (Tollefsen, 2002, p. 398). By this, Tollefsen means that they will not act against their best judgment, they will draw inductive inferences based on the available information, they will only believe things they think are true, and they will not hold conflicting beliefs (Tollefsen, 2002, p. 399). If the agent is not rational – or if we don’t think of it as rational – then it will remain entirely alien to us, utterly uninterpretable.

For groups, Tollefsen (2002) claims that their organisation provides a way of synthesising the various views of its members into a single, coherent perspective (p. 401). The hierarchy of a group itself enforces a means-end hierarchy of goals, where those lower down in the hierarchy perform actions to achieve goals that are only sub-goals for the goals of the members higher up the chain (Tollefsen, 2002, p. 401). So, Tollefsen (2002) argues that when we think of group agents as rational, we assume that they are structured in order to effectively pursue their goals (p. 401).

The issues with interpretationism are twofold. First, the intentional stance is far too broadly applicable (Keijzer, 2006, p. 1593). A thermostat can be conceived of as a rational, intentional agent that actively maintains and regulates the temperature in its environment because that is its goal. We can successfully predict what the thermostat will do in particular conditions given what we know of its “beliefs” about its environment. This is an issue in its own right, but it also points to the issue with making assumptions or assertions from an outside perspective rather than grounding normativity in the system that is itself supposed to have those norms. The interpretationist must *assume* that the agent they are dealing with shares their norms of rationality. On the organisational account, however, no such assumptions are necessary. Instead, the norms of the system are determined by the system itself on the basis of its organisational structure, which similarly belongs to the system. Even if interpretationism is

supposed to involve interpreting the agent based on what makes sense to it, it is still we who are doing the interpreting, and this creates an inherent gap between the norms and the agent who has the norms, meaning there is always going to be more room for error than there will be given an accurate ontological account that grounds norms in the agent itself.

The gap between the agent and interpreter leads to the second issue, which lies in the inability of interpretationism to account for how norms arise. That this is a problem is especially evident in the practices that follow from interpretationism. According to Tollefsen (2002), interpretationism applied to group agents is not just an ideal practice that we should strive for but is in fact commonplace (p. 402). Hence, we have laws that deal with groups rather than just some of their members without having required philosophers to convince lawmakers that these are justified. However, interpreting group agents in the same way that we interpret other people has resulted in largely ineffective laws. Slavery has still been the norm in chocolate production, for instance (Balch, 2021). Likewise, the environment is still being caused serious harm by large corporations, despite moral demands being made of companies and governments alike. We treat groups as if they can be dealt with in the same ways that people are dealt with. But groups cannot be shamed or reasoned with – they cannot be scared, or hurt, or worried. They must be dealt with differently, and to get at how exactly they must be dealt with, we need to understand the source of their normativity alongside what those norms are. There is an inherent barrier that can only be crossed once we understand them on their own terms, as the organisational account does, rather than merely interpreting them.

Interpretationism is, no doubt, often epistemically beneficial. In a pinch, we are likely to get very near the agent's own normative "reasons" for its behaviour by taking the intentional stance. However, it does not differentiate between actual agents and thermostats, which means that our assessment of a system's normativity will always be based on our own perspective and not on the agent's actual needs. This creates room for error and further implies that groups are

much more like us than they really are, which leads to issues, especially in the legal realm. The organisational account better positions us to address harmful group behaviours since we understand the basis of their normative actions as well as what those norms are, rather than just the latter.

Pauer-Studer, on the other hand, offers an ontological account of group agent normativity, drawing on the constitutivism of David Velleman and Christine Korsgaard. Pauer-Studer's account is positioned as a critique of List and Pettit (2011) who, says Pauer-Studer, fail to connect the *causal* and *normative* levels of agency (Pauer-Studer 2014, p. 1629). The causal level concerns whatever generates the actions of the agent, while the normative level is concerned with what the agent should do as the kind of agent that it is. Constitutivism, Pauer-Studer argues, connects these two levels. It tells us how the group's norms cause the group to act as it should, given the kind of agent it is. At its most basic, constitutivism is the view 'that the constitutive standards of agency provide a foundation for normativity' (Pauer-Studer, 2014, p. 1632). Put differently, the process of being constituted as an agent necessarily makes one a normative being. In order to persist as an agent, an animal necessarily needs food, water, and sleep at least. Hence, the continual constitution of themselves as an agent at all requires that they have normative commitments.

So far, the overlap between the organisational account and constitutivism is substantial. They differ, however, in the details – especially concerning the normative identity of an agent. The constitutive account requires that the agent have a coherent self-understanding (Pauer-Studer, 2014, p. 1635) and that its actions satisfy certain standards of intelligibility (p. 1636). In the case of group agents, their legal and organisational structures are defined when they are formed which, along with the group's specific aims, make up its normative identity (Pauer-Studer, 2014, p. 1635). This normative identity then determines whether the behaviours of the group allow for it to have a coherent self-understanding (Pauer-Studer, 2014, p. 1635).



According to Pauer-Studer (2014), this normative identity is generated in the first place by understanding the legal form of the group agent, i.e., whether it is a business, a political party, a university, and so on. The members of the group agent then determine the relevant decision making procedures and the structure of the group accordingly (p. 1635). Standards of intelligibility arise as a result of the fact that the group's actions can be judged in light of its structure, which is itself based on its initial constitution (Pauer-Studer, 2014, p. 1635).

An important point of difference between the organisational and constitutive accounts is in defining the identity of the group agent. According to the organisational account, the group identity is given by its organisational structure. If the group is structured to achieve political power using the mechanisms of the liberal democratic state, it is a political party. For the constitutivist, however, the agent is identified with all of its normative aims rather than its minimal, structurally-defined norms only (Pauer-Studer, 2014, p. 1636).

Consider the example that Pauer-Studer (2014) provides: if a corporate firm 'claims to uphold ethical standards such as respecting workers' rights and welfare' (p. 1636) but then undermines its workers' rights by outsourcing to countries with poor worker protection laws and exploiting those laws, then it undermines both its credibility and its 'coherence and intelligibility' (p. 1636). Hence, on the constitutivist reading of this case, it is a constitutive aim of the group that it behave morally with regard to its workers. The group's identity is constitutively tied with this moral aim. However, in the case that it fails to live up to this norm, it is unclear what is supposed to happen when the group stops being coherent and intelligible. Since it is a corporate firm, its failing to continue to act ethically would not obviously result in the firm's becoming unviable. It would certainly be unfortunate for its members, but the charge of incoherence and unintelligibility does not appear to have much weight in this case.

Using the organisational account, however, we can easily explain the situation. A corporate firm is aimed at making a profit and treating its workers ethically will reduce its

viability so long as being unethical is more profitable. As long as it does not reduce profit so much that the business starts to lose money, however, the moral norm can stay in place. But, whenever there is the need to increase its profit, one obvious way to do so will always be to outsource to cheaper workers. This is what I mean when I say that the adaptive demands of the system's viability always constrain and motivate the members and other norms present within the group. The ethical treatment of its employees is constrained insofar as the group must remain profitable in order to persist. It also motivates the unethical treatment of its workers since this will increase its adaptive viability. Rather than becoming incoherent and unintelligible, then, the shift in the group's activity is in fact an improvement in its viability. Given that companies do engage in harmful practices frequently, even if they were set up with good intentions, the organisational account rather than the constitutivist reading of this situation seems the more accurate of the two.

So, by identifying group agents with their structure, which is itself the source of the group's minimal, viability-related norms, the organisational account better addresses instances where groups do not behave according to previous, non-adaptive norms. The constitutive account instead places too much emphasis on all of the norms that the group has adopted at present, which naturally leads to the claim that the abandonment of prior moral norms leads to incoherence and unintelligibility. However, since the group would persist so long as it remains profitable, it is not clear what it means for the group to fail to be coherent and intelligible in this case. Despite their similarities, then, the organisational approach better accounts for the actual behaviour of group agents.

## **5 Conclusions**

I have argued here for the extension and application of the organisational account of normativity to group agents. According to the account, a group agent is at least an adaptive

system that is oriented toward its own self-maintenance given the demands of its present environment. Beyond this minimal level of normativity, groups are also capable of pursuing other norms given to them by their members. The activities of the members in their group-related roles, however, will always be constrained and motivated by the minimal normative requirements of their group, i.e., those norms related to its viability. It is possible that they can get the group agent to adopt norms that reduce its overall viability, so long as it is otherwise adaptively successful, but these norms will always be under pressure from the group's own normative tendency toward improving its viability. This pressure comes in countless forms, as I have shown above.

I will end the discussion by highlighting what the organisational account can tell us about what is required for changing the behaviour of collective systems.

The organisational account deals with both what a group's norms are and how they arise. It is this that really positions us to address unjust collective behaviours. Having an understanding of the norms of the agent will allow us to make an accurate assessment of its goals. By understanding how those norms arise, we can then deal with changing the group's behaviour by influencing the ways it can achieve its normative goals or by changing what those goals are. As I have argued, the viability of a corporate group agent is often improved by the poor treatment of its workers. So long as it is an option for them, businesses will be incentivised to outsource to less well-protected workers, to convince or trick their employees into doing unpaid overtime, or to enforce draconian measures to improve sales. These activities, though unjust, are adaptive for the group agent, since they take it further away from a state of failure. Since, as the organisational account indicates, these behaviours are not strictly the result of the personal choices of greedy or immoral individuals but are rather the appropriate responses by the group to its present environment, it will do us little good in influencing the behaviour of the group agent to think of the harms caused in terms of individual failings, as there is

sometimes a tendency to do. Instead, the organisational account emphasises the need for structural change or collective action to achieve long-term success. This could take the form of writing new laws, rewriting ineffective laws, taking part in worker strikes that threaten the ability of groups to function, or more radical action focused on changing the economic environment entirely.

It is important to emphasise that this result is one concerning the actions required for broader change. It is not a theory of the moral culpability of individuals. Pauer-Studer (2018) argues that an individual can be complicit in the actions of a group because they play a constitutive role in that group's activities, regardless of that person's own feelings about their professional identity (pp. 651-5). Holding someone morally culpable for harm is a separate issue to addressing the underlying mechanisms that result in that harm being done so consistently. Individual complicity in unjust collective actions is therefore perfectly compatible with the organisational account. The point is not that we should stop making moral judgments concerning the actions of individual people, but that if we wish to achieve long-term change, then we must also work proactively toward that change together. The organisational account, then, both accurately explains the normativity of group agents and provides us with a powerful tool for social and political theorising.

## References

- Balch, O. (2021). Mars, Nestlé and Hershey to face child slavery lawsuit in US. *The Guardian*, 13 February, viewed 13 March 2022. <<https://www.theguardian.com/global-development/2021/feb/12/mars-nestle-and-hershey-to-face-landmark-child-slavery-lawsuit-in-us>>
- Barandiaran, X. E., Di Paolo, E., & Rohde, M. (2009). Defining Agency: Individuality, normativity, asymmetry, and spatio-temporality in action. *Journal of Adaptive Behavior*, 17(5), 367-386.
- Barandiaran, X. E. & Egbert, M. (2013). Norm-Establishing and Norm-Following in Autonomous Agency. *Artificial Life*, 20(1), 5-28.
- Canguilhem, G. (1991). *The Normal and the Pathological*. C. R. Fawcett (trans.). Zone Books.
- Di Paolo, E., Buhrmann, T. & Barandiaran, X. E. (2017). *Sensorimotor Life: An Enactive Proposal*. Oxford University Press.  
<https://doi.org/10.1093/acprof:oso/9780198786849.001.0001>
- Jonas, H. (1966). *The Phenomenon of Life: Toward a Philosophical Biology*. Northwestern University Press.
- Keijzer, F. (2006). Differentiating Animality from Agency: Towards a Foundation for Cognition. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 28, 1593-1598.
- List, C. & Pettit, P. (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford University Press.  
<https://doi.org/10.1093/acprof:oso/9780199591565.001.0001>
- List, C. (2021). Group Agency and Artificial Intelligence. *Philosophy and Technology*, 1-30.  
<https://doi.org/10.1007/s13347-021-00454-7>

- Pauer-Studer, H. (2014). A Constitutive Account of Group Agency. *Erkenntnis*, 79(9), 1623-1639. <https://doi.org/10.1007/s10670-014-9632-y>
- Rand, S. (2011). Organism, normativity, plasticity: Canguilhem, Kant, Malabou. *Continental Philosophy Review*, 44(1), 341-357. <https://doi.org/10.1007/s11007-011-9196-3>
- Rovane, C. (2019). Is group agency a social phenomenon?. *Synthese*, 196(12), 4869-4898. <https://doi.org/10.1007/s11229-017-1384-1>
- Stapleton, M. & Froese, T. (2015). Is Collective Agency a Coherent Idea? Considerations From the Enactive Theory of Agency. In C. Misselhorn (ed.), *Collective Agency and Cooperation in Natural and Artificial Systems: Explanation, Implementation and Simulation* (pp. 219-236). Springer. <https://doi.org/10.1007/978-3-319-15515-9>
- Tollefsen, D. (2002). Organizations as True Believers. *Journal of Social Philosophy*, 33(3), pp. 395-410. <https://doi.org/10.1111/0047-2786.00149>
- Tuomela, R. (2013). *Social Ontology: Collective Intentionality and Group Agents*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199978267.001.0001>