



ARTICLE

# Bratman on institutional agency

Jessica Brown 

Arché Philosophical Research Centre, University of St Andrews, St Andrews, UK

Email: [jab30@st-andrews.ac.uk](mailto:jab30@st-andrews.ac.uk)

(Received 28 February 2024; revised 16 May 2024; accepted 12 September 2024)

## Abstract

In his recent book, *Shared and Institutional Agency*, Bratman attempts to defend realism about institutional agency while appealing only to ontologically modest foundations. Here I argue that this ontologically modest foundation leaves Bratman unable to provide plausible accounts of institutional evidence, institutional belief, and the reasons for which institutions believe and act. Given that these phenomena are key to our moral and epistemic evaluation of institutions and their actions, this is a serious failing. Instead, we should defend a more robust realism about institutional agency.

**Keywords:** Institutional agency; institutional evidence; institutional belief; institutional reasons for action and belief; Bratman

## 1. Introduction

In one of the most important recent books on institutions, Bratman attempts to defend realism about institutional agency while appealing only to ontologically modest foundations. In particular, he attempts to construct institutional agency out of facts about individual agents together with his previously defended account of small-scale shared agency. Here I argue that this construction fails to provide a suitable account of institutional agency. In particular, I argue that appealing to this ontologically modest foundation leaves Bratman unable to provide plausible accounts of institutional evidence, institutional belief, and the reasons for which institutions believe and act. Given that these phenomena are key to our moral and epistemic evaluation of institutions, this is a serious failing.

## 2. Bratman's ontologically modest account of institutional agency

Bratman (2022) argues that institutions are genuine agents with intentions and beliefs. Further, he argues that institutional agency can be constructed out of facts about individual planning agents and the kind of shared intentions he has previously argued are central to small-scale shared intentional activities. He contrasts his view with two rival approaches – ‘strongly individualistic planning approaches’ and ‘parallel planning approaches’ (xvii). In contrast to his realism about institutional agency, strongly individualistic planning approaches attempt to characterise institutions by appealing to merely personal plans including the kind of shared intentions characteristic of small-scale

© The Author(s), 2025. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

intentional activity, e.g. singing a duet or moving a piano upstairs. However, strongly individualistic planning approaches face well-known problems. For instance, unlike shared intentions, institutional intentions are robust to changes in participants, can persist despite ‘alienated’ participants who don’t intend the relevant outcome or even know what it is; and may involve members who are not directly in touch with others and so cannot have the kind of *de re* intentions and attitudes to each other Bratman thinks characteristic of small-scale shared intentional activity (e.g. Shapiro 2014, Habgood-Cooté 2020, Ritchie 2020). Given these problems, some instead endorse a ‘parallel planning approach’ which attributes planning attitudes and activities directly at the level of the institution without constructing them out of an interpersonal web of individual and shared plans and planning. Like the parallel planning approach, Bratman is realist about institutional agency but he hopes to defend it on more ontologically modest foundations, showing how it can be constructed out of individual and shared plans.

Key to Bratman’s construction is the notion of a ‘social procedural rule’ of an institution. Inspired by Hart’s account of social rules, he treats social procedural rules as having two aspects: an external aspect consisting of a social behavioural regularity and an internal aspect that consists in the relevant standard being endorsed. The relevant regularity is a regularity in following certain institutional procedures for resolving problems. And he aims to explain the internal aspect by appealing to his notion of shared intention, while acknowledging that it will need to be tweaked to play this role given the differences detailed above between institutional intention and shared intention (Chapter 4). Once he has characterised social procedural rules, Bratman argues that their outputs constitute mental states of the institution, such as intentions, acceptances, and beliefs.<sup>1</sup> In particular, if a social procedural rule favours an action, then the institution counts as intending that action. And if a social procedural rule favours a proposition, then the institution counts as accepting that proposition. However, Bratman admits that institutional acceptances will not in general be beliefs for they may be sensitive to a variety of practical considerations in a way characteristic of acceptance rather than belief. For instance, whether an oil company accepts some proposition about climate change might be affected by consideration of its profits. However, where institutional acceptances are tightly tied to the evidence, Bratman treats them as constituting institutional beliefs (139).

Although Bratman recognises that to provide an account of institutional belief, he needs an account of institutional evidence, he provides no such account. And I’ll argue in the next section that there is no obvious way for him to provide a satisfactory account of institutional evidence.

### 3. Bratman and institutional evidence

The basic elements of Bratman’s construction of institutional agency are facts about individuals plus social procedural rules. Let’s consider whether either could provide

---

<sup>1</sup>More precisely, it is “crystallised” outputs that constitute institutional mental states. A crystallised output is a temporarily-extended social-psychological web that favours the relevant output (104). For instance, if the output favours accepting *p* then, if all goes well, *p* will appear as a premise in the further reasoning of the institution or its individual members. A crystallised output counts as the output of a certain institution if and only if 1) it is consistent with other outputs or licensed by an authorised adjudication rule; 2) tends to shape relevant downstream thought and action in relevant ways where that entails 3) limited functionally relevant non-compensated cognitive-divergence among members in belief about means or relevant factual claims. (126). This further precisification doesn’t affect the argument to come and so I omit discussion of it. For instance, even acceptances which are not sensitive to evidence and so are not plausibly beliefs might appear as premises in the further reasoning of an institution.

Bratman with an adequate account of institutional evidence. There are two broad ways in which one could attempt to use social procedural rules to provide an account of institutional evidence: one might treat an institution's evidence either as the outputs of, or the inputs to, its social procedural rules. Treating the outputs of social procedural rules as an institution's evidence is very unpromising. Action-oriented outputs are the wrong kind of thing to provide an account of institutional evidence. And, as we have already seen, acceptance-oriented outputs are typically shaped by pragmatic concerns such as the bottom line in a way which is incompatible with their serving as evidence. Bratman's account of institutional beliefs as constituted by those institutional acceptances which are tightly tied to evidence presupposes an account of institutional evidence, rather than providing one.

Instead, it might be suggested that an institution's evidence can be identified with the inputs to its social procedural rules, or perhaps what the institution's processes deem acceptable as input to its social procedural rules. However, identifying an institution's evidence with what it in fact inputs, or deems acceptable as input, to its social procedural rules is both too inclusive and not inclusive enough. It's not inclusive enough because an institution might illegitimately rule out some evidence as acceptable input to its social procedural rules. For instance, an oil company might not treat as acceptable input to its social procedural rules well-attested facts about climate change. Furthermore, it's too inclusive since an institution might input certain propositions into its social procedural rules even though they don't constitute evidence.<sup>2</sup> For instance, an oil company might treat as input to its social procedural rules fringe views on climate change which have no scientific backing because of the company's desire to continue oil production.

It seems, then, that appeal to neither the inputs nor outputs of social procedural rules will provide a satisfactory account of institutional evidence. Setting aside social procedural rules, the other basic element of Bratman's construction are facts about individual agents. So, Bratman might suggest that the evidence of an institution is a function of the evidence of its members. Thus, it may be suggested that *p* is part of the evidence of an institution if and only if it is part of the evidence of some/most/all members of the institution.<sup>3</sup> Of the various possible functions, we can quickly set aside the suggestion that *p* is part of the evidence of an institution if and only if it is part of the evidence of *every* member of the institution. For, surely, e.g. a government's evidence could include that, say, a foreign agent has hacked its taxation system even if one member of the government doesn't know this because they fell asleep during the relevant meeting. So, that leaves us with the some/most readings. However, as I have argued elsewhere, both of these face serious problems: 1) on a non-factive account of evidence, they will likely result in institutional evidence being inconsistent; 2) they yield too inclusive accounts of evidence; and 3) they violate a plausible and widely endorsed doxastic constraint on evidence (Brown 2022a). Let me sketch the problems, focusing on the following 'pooled' account:

---

<sup>2</sup>One might hope to ameliorate the too-inclusive worry by suggesting that an institution's evidence consists in those inputs to its SPR which meet some further epistemic condition. It wouldn't be sufficient to add that the inputs be true for, intuitively, true inputs don't always constitute part of the institution's evidence (e.g. if they are the result of wishful thinking which just happens to be true). Further, since Bratman lacks an account of institutional justified belief/knowledge he cannot add the condition that the inputs are justifiably believed/known by the institution. Alternatively, it might be suggested that the inputs need themselves to be the product of a further SPR which is reliable. But this wouldn't avoid the failure to be inclusive enough and, in any case, would be philosophically controversial (e.g. it would be a non-factive account of institutional evidence and would allow that reliable inference expands an institution's evidence).

<sup>3</sup>Functions could also differ in which are the relevant members to consider. I leave that issue aside.

*Pooled*: p is part of the evidence of an institution if and only if it is part of the evidence of some member of the institution.

The first problem is that if *Pooled* is combined with a non-factive account of evidence, then it will likely result in institutions having inconsistent sets of propositions as evidence.<sup>4</sup> For instance, assume that it is sufficient for a proposition to be part of a subject's evidence that she justifiably believes it. On such a non-factive approach, it's possible that the different members of a single institution have justified beliefs in inconsistent propositions. For instance, one member might have p as part of her evidence, but another member might have not-p as part of her evidence. Thus, the institution's evidence would include both p and not-p.

A second problem facing *Pooled* is that it yields counterintuitive results about institutional evidence. For instance, suppose that a male member of a department, X, has been sexually harassing a female colleague, Y. Given their involvement, both X and Y have evidence that such sexual harassment has taken place (indeed, they both know it has!). On *Pooled*, if a member of an institution has a proposition, p, as part of her evidence, then so does the institution. Thus, on *Pooled*, the department's evidence includes the proposition that X has sexually harassed Y. But that may be counterintuitive depending on how we fill out the details of the case (Brown 2022a 167–168). Suppose, for example, that no member of the department besides X and Y has any evidence of the sexual harassment since neither X nor Y reveal the sexual harassment (perhaps, X doesn't want to be sanctioned and Y has decided to deal with the incident in her own way). This is despite the fact that the department has excellent systems for dealing with allegations of sexual harassment; indeed, we may imagine the department systems are just as we would want them to be. In such circumstances, it seems intuitive that the department has no evidence that X has sexually harassed Y. This is backed up when we consider the department's actions. Given that the department is committed to gender equality and eliminating sexual harassment, it would be puzzling if the department had evidence of harassment yet did nothing. Indeed, the department might well be morally blameworthy for failing to act if it were to have such evidence. Thus, *Pooled* seems to be overly generous when it ascribes the evidence of every member of an institution to the institution itself.<sup>5</sup>

Even setting aside the first two concerns, *Pooled* faces the further objection that it is in tension with the popular doxastic constraint on evidence according to which p is part of one's evidence only if one believes that p. To illustrate the problem for *Pooled*, assume for the sake of argument that one's evidence is one's knowledge. Further, assume that just one of 20 operative members, M1, of an institution knows that p. By the identification of evidence and knowledge, M1's evidence includes that p. According to *Pooled*, a proposition p is part of an institution's evidence if and only if it is part of the evidence of some member of the institution. As a result, the institution's evidence

<sup>4</sup>This problem would not arise for a factive account of evidence.

<sup>5</sup>Notice that the problem still arises even if we strengthen the conditions on institutional evidence such that p is part of a group's evidence if and only if 1) p is part of the evidence of most members of the group; and 2) 1) is common knowledge (Brown 2022a: 168). Suppose that it's part of the evidence of most members that sexual harassment is occurring because there are multiple victims of the harasser, and also suppose that this is common knowledge. That is compatible with no one reporting the harassment to the group's processes. But, if none of the victims come forward to report the harassment to the group's processes, it may be plausible that the group's evidence does not include the fact that sexual harassment is going on, and so the group would be morally blameless for failing to sanction the perpetrator. Of course, that's compatible with various members of the group being morally blameworthy (most obviously the perpetrator) and the group being legally liable for the harassment (for legal and moral standards may differ).

includes *p*. However, on neither a summative nor a non-summative account of institutional knowledge, does it follow that the institution knows that *p*. A summative account of institutional knowledge would plausibly hold that an institution knows that *p* only if most of the members of the institution know that *p*, or perhaps most of the operative members do.<sup>6</sup> But, this condition is not met in our example. On a non-summative account of institutional knowledge, whether an institution knows that *p* is not a function of whether its members know that *p*: that all or most members know that *p* does not entail that the institution knows that *p*; and that the institution knows that *p* does not entail that any of its members know that *p*. Thus, from the fact that one member of the institution knows that *p* nothing follows about whether the institution knows that *p* or not. Thus, the pooled conception of institutional evidence is in tension with the doxastic constraint on evidence: it allows that a proposition *p* may be part of an institution's evidence even though the institution doesn't hold the relevant doxastic relationship to *p*.<sup>7</sup>

In conclusion, neither of the main two elements of Bratman's modest ontological foundation – facts about individuals or social procedural rules – provide a plausible account of institutional evidence. In the next section, I turn to consider what account Bratman can offer of the reasons for which institutions act and believe.

#### 4. Reasons for belief and action

To enable us to epistemically and morally evaluate institutions, Bratman needs an account of the reasons for which institutions act and believe. For, whether an agent's action is praiseworthy or blameworthy, and whether her belief is justified or not is affected by the reason for which she acts or believes. For instance, two siblings might both help their aged aunt but for different reasons: one does it because she needs help but the other does it to increase the chance of a large inheritance. Only the former is praiseworthy for helping her aunt. Likewise, the evaluation of the beliefs of an agent depends on their basis. For instance, contrast two detectives who both believe that *X* was responsible for the murder of a politician, where one believes that on the basis of supporting evidence, but the other believes that out of wishful thinking. Only the former justifiably believes that *X* murdered the politician. Similar points apply to our evaluation of the actions and beliefs of institutions. We might differently evaluate the actions of two companies who each implement a green policy but do so for different reasons: to help the environment or to increase profits by attracting the green pound. Only the former is praiseworthy for helping the environment. Likewise, we might differently assess the beliefs of two governments who each believe that it's safe to terminate social distancing measures but do so for different reasons: the evidence or because it will increase votes in the upcoming election. Only the former has a justified belief.

<sup>6</sup>It's not plausibly necessary for an institution to know that *p* that every member does (e.g. the government can know that inflation is forecast to fall even if one member doesn't know that because she missed the relevant meeting). Nor is it plausibly sufficient for an institution to know that *p* that at least one member knows that *p*. For instance, it's not sufficient for Prof *X*'s department to know that Prof *X* has signed a contract to take a job elsewhere that just one member of the department does. For if Prof *X* keeps this knowledge to herself, the department plausibly does not know she's signed the contract.

<sup>7</sup>Notice that it's no help to appeal to the idea that the doxastic constraint only requires a potential doxastic relationship, say being in a position to know. For, the relevant notion of being in a position to know is usually understood to exclude the idea that one could come to know that *p* by further investigation. For, otherwise, it would have the implausible result that, say, a person's evidence includes everything she could come to know by reading Wikipedia. So understood, an institution is not automatically in a position to know everything which its members know.

Despite the importance of the reasons for which institutions believe and act, Bratman says very little about what it is for an institution to believe or act for one reason rather than another. Instead, in discussing Davidson's constraints on mental states, he stresses how, on his view, an institution's having one mental state doesn't require that it have a dense holistic web of other mental states (148). In particular, he explains how an institution can have an intention without there being an institutional ground for that intention. In his example, he supposes that a medical supply company intends to send medical supplies to a certain country, even though there is no institutional ground or reason for that intention. As he imagines it, different individuals within the institution support sending the supplies for different reasons (that it will reduce human suffering, because of promises made to donors, because of legal requirements, to support the relevant local government etc). Further, he says, 'there may be no need for the institution itself to take a stand about which, if any, of these potential rationales is to be incorporated within its institutional commitments'. Thus, at least for institutions, he rejects Davidson's idea that attributing a certain mental state to an agent requires the attribution of a dense pattern of other mental states that coheres with and makes sense of the relevant target attitude. However, even if Bratman is right that Davidson's holism about mental states doesn't apply to institutions, it remains the case that in order to evaluate institutional beliefs and intentions, we need to make sense of institutional reasons for belief and intention. Indeed, as Bratman recognises, sensitivity to evidence seems central to the very notion of belief and distinguishes belief from mere acceptance.

How, then, might Bratman attempt to add reasons for belief and intention into his picture? In discussing intention, he suggests that the social procedural rule of an institution could issue in outputs that involve further aims, what he calls 'induced teleology' (109). For instance, the social procedural rule of the medical supply company might issue a more detailed institutional output: send supplies in order to alleviate human suffering. And he says that when the social-rule-induced output has this further teleological structure, 'the activity that is guided by this output will be constrained to be responsive to that built-in further aim' (110). In this way, Bratman might hope that his framework can provide an account of when an intention is adopted for some particular end and thus, relatedly, when the institution acts for that end. For example, if the medical supply company acts on its intention to send medical supplies in order to alleviate human suffering then it would count as sending the supplies for the reason that it alleviates human suffering. As a result, the company would be praiseworthy for sending medical supplies. By contrast, if the relevant social procedural rule instead had output the intention to send the supplies in order to increase profits, then the institution wouldn't have been praiseworthy for sending the supplies.

However, I'll argue that such cases of induced teleology are neither necessary nor sufficient for an institution to act for a reason. To see that induced teleology isn't necessary, notice first that, as Bratman agrees, in many cases, social procedural rules may simply specify an action without specifying any end of the action. Despite that, it may be intuitively plausible that the institution does act for a particular reason and is accordingly praiseworthy or blameworthy. For instance, consider two different companies both of which instigate the same green policy, even though, for neither, does the relevant social procedural rule specify the further end of the green policy. Even so, it may be intuitively plausible that just one of these companies instigates the green policy for the reasons that make it the right thing to do and is thus praiseworthy. Perhaps, one is a profit-driven company which always puts profit before any other consideration, whereas the other has a track record of accepting lower returns on investment for environmental goals. In these circumstances, it is intuitively plausible that only the second company instigates the green policy for the reasons that make it

right and is praiseworthy for doing so. Thus, it is not necessary for an institution to act for a certain end that it has a social procedural rule which specifies that the action is undertaken for that end.

Furthermore, that an institution has a social procedural rule which specifies that a certain action is undertaken for a certain end doesn't seem sufficient for the institution to act for that end. For, pragmatic considerations may lead its social procedural rule to specify that it is undertaking an action for a certain end, when in fact it is undertaking it for a very different end. We can illustrate this by a variant of the example of a profit-driven oil company instigating a certain green policy. We may suppose that rather than failing to specify the end for which it implements the green policy, its social procedural rule outputs that the green policy is undertaken for the sake of the environment. For instance, perhaps one of its institutional outputs is to reduce disturbance in a river delta in order to safeguard birdlife. Nonetheless, that leaves it open what is the real reason for which its social procedural rule outputs the green policy. For pragmatic reasons, it may suit the company to have a social procedural rule whose output specifies that it is adopting the green policy for the sake of the environment, even though in fact it only adopts it for the sake of increasing profits. For, the company may be aware that customers and investors are increasingly committed to environmental ends, and so it's in the company's interest to appear to be concerned with the environment, even though it's only concerned with profit.

Notice that the sufficiency problem is not solved by Bratman's idea that when an institutional output has this induced teleological structure, the activity guided by this output will be constrained to be responsive to the specified end. For, even a profit-driven company might have reason to ensure that the relevant action is responsive to the specified environmental end. For example, reconsider the profit-driven oil company with a social procedural rule whose output is to reduce activity in a river delta to safeguard birdlife. It's compatible with this company being driven fundamentally by profit that the relevant output – here reducing activity in the river delta – is responsive to the specified end – here to safeguard birdlife. For, attracting the green pound may require some policies which actually benefit the environment if, for instance, green investors monitor the upshot of the company's policies. So, it's in its financial interest to ensure that the river delta policy actually benefits birds.

We have now seen that Bratman's appeal to induced teleology is neither necessary nor sufficient for providing an account of the reasons for which institutions act. For similar reasons, it would be hard to employ his account of induced teleology to provide an account of the reasons for which institutions believe. Just as the action-oriented outputs of social procedural rules could include a specification of the ends of the action, so the acceptance-oriented outputs of social procedural rules could include a specification of the reasons for the acceptance. Thus, instead of an acceptance-oriented social procedural rule merely outputting, say, inflation will fall, it could instead output: inflation will fall because interest rates are rising. Thus, Bratman could suggest that when an institution's social procedural rule outputs: accept *p* because *q*, then the reason for which the institution accepts *p* is *q*.

But this account of the reasons for which institutions accept certain claims seems subject to the same objections that faced the induced teleological account of the reasons for which companies act. It is neither necessary nor sufficient for an institution to accept *p* for the reason that *q* that the acceptance is the output of a social procedural rule of the form: accept *p* because *q*. It's not necessary since an institution's social procedural rule might fail to specify any reason for its acceptance of *p* yet, intuitively, it accepts *p* for reason *q*. For instance, a profit-oriented oil company might have a social procedural rule which outputs: accept that there is no human-induced climate change without specifying



the reason for this acceptance where, intuitively, the reason for the acceptance is that it doesn't want to change its profitable business model. Further, it's not sufficient since an institution's social procedural rule might issue the output: accept  $p$  because  $q$ , even though, intuitively, the institution accepts  $p$  for some different reason,  $r$ . For instance, the social procedural rule of a government might issue the complex output: accept there is no anthropogenic climate change because so-called evidence for such climate change has been fabricated by left-wing agitators. But, we can imagine circumstances in which, intuitively, the reason for the government's denial of anthropogenic climate change is not that apparent evidence has been fabricated by left-wing agitators. Instead, it may be intuitive that both the government's denial of anthropogenic climate change and the reason it offers for its denial are both driven by fear of voter backlash. Thus, just as in the case of action, an institution's social procedural rule might specify the reason for its acceptance for pragmatic considerations, and so needn't reflect the real reason for which the acceptance was adopted.<sup>8</sup>

We've now seen that Bratman's appeal to induced teleology fails to offer a plausible account of the reasons for which an institution acts or believes. Given this failure, it might be suggested that Bratman should instead adopt a summative account of the reason(s) for which institutions believe or act. For this would be in keeping with his ontologically modest approach of constructing institutional agency out of individual planning agents and their shared intentions. Of course, it's too demanding to hold that an institution performs  $\alpha$  for reason  $r$  only if every member of that institution supports the institution's performing  $\alpha$  for  $r$ . Surely, for example, a government might intend to lower taxes to stimulate the economy even if one minister is motivated only by the desire to reduce taxes on his personal company. So it might instead be suggested that an institution performs  $\alpha$  for the reason that  $r$  if and only if enough of the relevant members of the institution support  $\alpha$  for  $r$ . Similarly, it might be suggested that an institution believes that  $p$  for the reason that  $r$  if and only if enough of the relevant members of the institution believe that  $p$  for  $r$ .

A first problem with the proposed threshold account concerns what counts as 'enough' of the relevant members and what determines that level. Presumably, it's a high threshold. But, wherever the threshold is placed – say above 70% – why should it be 70%, rather than 71%, 72%, or 73%? It's hard to see how the threshold is determined in a non-arbitrary fashion. But if the threshold is determined in an arbitrary fashion, that's deeply unappealing. For, after all, surely it is not arbitrary whether an institution acts for one reason rather than another and so whether it is morally blameworthy for its actions? (Brown 2021).<sup>9</sup>

A second problem concerns cases in which there is such a diversity in the reasons of members that, on the proposed threshold view, the institution doesn't count as acting or

---

<sup>8</sup>It wouldn't help Bratman to instead focus on beliefs, rather than acceptances. First, to distinguish beliefs and acceptances he needs the notion of institutional evidence but, as we saw earlier, he doesn't have an adequate account of institutional evidence. Second, it would be problematic to suggest that an institution believes that  $p$  for  $r$  if and only if it believes that  $r$  is a good reason to believe that  $p$ . For, this amounts to a doxastic account of basing which faces well-known objections. Having a meta-belief about the reason for which one has a belief isn't necessary for that belief to be based on that reason (consider subjects without the concept of belief). Nor is it sufficient since the subject could believe that  $r$  is a good reason to believe that  $p$  but not believe it for that reason. For instance, someone might believe that the evidence is a good reason to believe that their partner is faithful but, instead, believe it from wishful thinking.

<sup>9</sup>It's not obvious that appeal to context helps to avoid the problem. For instance, that the institution is a government and that the intention concerns important matters of public interest (tax rates) doesn't obviously determine what proportion of members of the government supporting lower taxes for  $r$  is necessary and sufficient for the government to lower taxes for  $r$ .



believing for any particular reason. Thus, the institution doesn't count as positively evaluable for the relevant belief or action. Nevertheless, in some such cases, the institution is intuitively positively evaluable for the relevant belief or action. We can illustrate the point in the case of belief. The different members of an institution may each believe a proposition,  $p$ , for different reasons, each of which justifies the belief that  $p$ . But if there is sufficient diversity amongst these reasons, then, on the proposed threshold view, there is no reason for which the institution believes that  $p$ . But, given that each member believes that  $p$  on the basis of a justifying reason, it may seem plausible that the institution justifiably believes that  $p$ . (Brown 2022b).

For instance, Lackey (2021) provides an example in which a group of 60 museum guards each justifiably believes that some guard is planning an inside theft of a painting from a museum, but do so on different but compatible bases. Each of guards M1-M20 justifiably believes that guard Albert is planning a theft and so deduces that some guard is planning an inside theft; none of the other guards believe that Albert is planning a theft. Each of a second group of guards M21-M40 justifiably believes that guard Bernard is planning a theft and deduces that some guard is planning an inside theft; none of the other guards believe that Bernard is planning a theft. Each of M41-M60 justifiably believes that guard Cecil is planning a theft and so deduces that some guard is planning an inside theft; none of the other guards believe that Cecil is planning a theft. Intuitively, the group of guards justifiably believes that some museum guard is planning an inside theft, even though there is no single reason for which most of the guards believe that (e.g. Goldman 2014, Lackey 2021, Silva 2019).

Thus, the proposed threshold account fails to provide a plausible account of cases in which members of an institution have a diversity of reasons for the relevant belief or action. Instead of defending the threshold account, a summativist might suggest that the reason for which the institution acts or believes is a conjunction of the reasons for which its members act or believe. For instance, the reason for which the group of museum guards believes that some guard is planning an inside theft of a painting from the museum is the following conjunction (Albert is planning a theft and Bernard is planning a theft and Cecil is planning a theft). This conjunctive result would be produced by the following account of an institution's reasons for belief/action:

*Conjunction*:  $r$  is a reason for which an institution believes that  $p$ /performs  $\alpha$  if and only if  $r$  is a reason for which some member believes that  $p$ /plays their part in the institution's  $\alpha$ -ing.<sup>10</sup>

*Conjunction* seems to get the intuitively correct results in *Museum Guard*.

Although *Conjunction* seems to be initially appealing, it has problematic consequences. To set up the problems, notice that *Conjunction* needs to be conjoined with an appropriate account of institutional reasons. For,  $r$  can be a reason for which an agent believes/acts only if  $r$  is one of that agent's reasons. It would be natural to combine *Conjunction* with a pooled account of institutional reasons:

*Pooled reasons*:  $r$  is one of an institution's reasons if and only if it is one of the reasons of some member of the institution.

Indeed, we can show that *Conjunction* entails *Pooled Reasons*. To see this, first suppose that  $r$  is the reason for which an institution performs  $\alpha$ /believes that  $p$ . It follows that  $r$  is

<sup>10</sup>This formulation reflects the fact that when an institution performs a certain action,  $\alpha$ , it's possible that no member individually performs that very action but instead plays a suitable part in it.

one of that institution's reasons. Now suppose that the necessity direction of *Pooled Reasons* fails, i.e. it's possible that *r* is one of an institution's reasons and yet it is not one of the reasons of some member of that institution. So, let's suppose that *r* does indeed fail to be a reason of any member of the institution. It follows that *r* is not a reason for which any member of the institution believes or acts. Thus, the necessity direction of *Conjunction* fails: *r* is a reason for which an institution acts or believes but is not a reason for which any of its members act or believe.

Now, let's see how the falsity of the sufficiency direction of *Pooled Reasons* would undermine the sufficiency direction of *Conjunction*. Start by assuming that *r* is a reason for which some member of the institution acts or believes. Thus, *r* is a reason of that member. If the sufficiency direction of the *Pooled Reasons* fails, then it is nonetheless possible that *r* is not a reason of the institution. But in that case, *r* cannot be a reason for which the institution acts or believes. Thus, the sufficiency direction of *Conjunction* fails. There is a case in which *r* is a reason for which some member believes that *p* or plays their part in  $\alpha$ , yet *r* is not a reason for which the institution believes that *p* or performs  $\alpha$ .

We have now seen that *Conjunction* entails *Pooled Reasons*. But as we can see from the earlier discussion of evidential reasons for belief (section 3), *Pooled Reasons* faces a range of serious objections. In particular, it's overly inclusive, incompatible with the doxastic constraint, and if combined with a non-factive account of reasons, results in institutions having inconsistent sets of reasons. Given the problems earlier raised for the threshold account, it seems that there is no plausible summative account of the reasons for which groups act and believe available to Bratman.

I've now argued that the main resources Bratman allows himself – facts about individuals and social procedural rules – are insufficient to provide a plausible account of the reasons for which institutions believe and act. Together with the earlier result that these resources are insufficient to provide a plausible account of institutional evidence, which in turn is required by Bratman's account of institutional belief, we see that his resources are inadequate for three central phenomena that we use to assess institutions – institutional belief, institutional evidence, and the reasons for which institutions believe and act.

## 5. Conclusion

I've been evaluating Bratman's attempt to provide a realist account of institutional agency while appealing only to ontologically modest foundations, specifically facts about individual agents and the kinds of shared intention he thinks are characteristic of small-scale shared intentional activities such as singing a duet or painting a fence together. Bratman appeals to shared intention to help explain his key notion of a social procedural rule – a regularity in how institutions problem-solve which is endorsed by the institution for the institution. In his approach, the intentions and beliefs of institutions are understood as certain kinds of outputs of social procedural rules, respectively action-focused and acceptance-focused outputs. However, I've argued that his account of institutional agency faces serious problems. His account of institutional belief presupposes an account of institutional evidence. But Bratman fails to provide any such account and I've argued that the resources he allows himself don't enable him to provide a plausible account. Further, I've argued that those same resources don't enable him to offer a plausible account of the reasons for which institutions believe or act. This is a serious failing for these phenomena – institutional belief, institutional evidence, and the reasons for which institutions act and believe – are key to our evaluation of institutions, and whether their beliefs and actions are justified or not. Bratman was motivated to endorse his realist yet ontologically modest approach to institutional

agency by the well-known problems facing strongly individualistic planning approaches which attempt to characterise institutions merely by appealing to personal plans. In the light of the difficulties facing strongly individualistic planning approaches and Bratman's own approach, I suggest that one ought to explore what Bratman calls a parallel planning approach which attributes planning attitudes and activities directly at the level of the institution without constructing them out of an interpersonal web of individual and shared plans.

**Competing interests.** The author declares none.

## References

- Bratman M.** (2022). *Shared and Institutional Agency*. Oxford: OUP.
- Brown J.** (2022a). 'Group Evidence.' *Philosophical Issues* 32(1), 164–79. DOI: [10.1111/phis.12220](https://doi.org/10.1111/phis.12220)
- Brown J.** (2022b). "'Group Belief for a Reason". Presidential Address of the Joint Session of the Aristotelian and Mind Societies 2022.' *Proceedings of the Aristotelian Society, Supplementary Volume* 96(1), 1–22. <https://doi.org/10.1093/arisup/akac004>
- Brown J.** (2021). 'Group Motivation'. *Nous* 56(2), 494–510. doi: [10.1111/nous.12366](https://doi.org/10.1111/nous.12366)
- Goldman A.** (2014). 'Social Process Reliabilism: Solving Justification Problems in Collective Epistemology'. In J. Lackey (ed.), *Essays in Collective Epistemology*, pp. 11–41. Oxford: Oxford University Press.
- Habgood-Coote J.** (2020). 'Group Knowledge, Questions, and the Division of Epistemic Labour.' *Ergo* 6(33), 1–42.
- Lackey J.** (2021). *The Epistemology of Groups*. Oxford: Oxford University Press.
- Ritchie K.** (2020). 'Minimal Cooperation and Group Roles.' In A. Fiebich (ed.), *Minimal Cooperation and Shared Agency*, pp. 93–109. Dordrecht: Springer.
- Shapiro S.** (2014). 'Massively Shared Agency.' In M. Vargas and G. Yaffe (eds), *Rational and Social Agency: the Philosophy of Michael Bratman*, pp. 257–293. Oxford: Oxford University Press.
- Silva Jr. P.** (2019). 'Justified Group Belief is Evidentially Responsible Group Belief.' *Episteme* 16, 262–81.

**Jessica Brown** is Professor of Philosophy in the Department of Philosophy at St Andrews University. Her books include *Groups as Epistemic and Moral Agents* (OUP, 2024); *Fallibilism: Evidence and Knowledge* (OUP, 2018); and, *Anti-Individualism and Knowledge* (MIT, 2004).