Replies to Hawley, Mikkola, and Hindriks

Brian Epstein

Replies to Hawley, Mikkola, and Hindriks

Brian Epstein

Department of Philosophy, Tufts University, Medford, MA, USA

I am very grateful for these insightful and challenging comments. I will begin with the comments on anchoring and grounding by Hawley and Mikkola, and then discuss group agency and Hindriks's defense of the status account.

Initial clarifications

Both Hawley and Mikkola usefully introduce parts of the first half of *The Ant Trap (Epstein 2015)*. Hawley starts with a compelling discussion of dependence relations in social metaphysics, and Mikkola endorses my arguments against 'anthropocentrism' in social ontology and explains the grounding–anchoring model. To both of these introductions, however, I should add a couple of clarifications and corrections.

As Hawley correctly says, one aim of the first half of *The Ant Trap* is to articulate the 'Standard Model' of social ontology, as developed by philosophers such as David Hume, H.L.A. Hart, and John Searle, among many others. Searle's 'collective acceptance' theory is perhaps the best-known contemporary variant of this model. But it may be worth adding that this is just one of several goals of the first half of the book. More generally, the first half aims to develop a new and unified model for social ontology to replace two different and competing traditions for analyzing the nature of the social world. One tradition is largely associated with ontological individualists such as J.S. Mill and Karl Popper. The other is the tradition to which the 'Standard Model' belongs. I argue that these two traditions should be understood as theorizing about distinct aspects of how the social world is built. I argue that ontological individualism ought to be understood as a theory of how social facts are *grounded*, and the 'Standard Model' as a theory of how the social world is *anchored*. A better theory of

the social world addresses both anchors and grounds – and addresses them in a more sophisticated way than these prevailing models do.

A second important clarification – or rather, correction – is that I do not endorse the 'Standard Model', and certainly do not endorse a collective acceptance theory (which is Searle's particular version of it). At one point, Mikkola suggests that I do, at least for the case of money, but this is not my view. In explaining both the 'Standard Model' and the grounding–anchoring model, I do use collective acceptance to illustrate a theory of anchoring. However, I chose Searle's treatment of money as an example – both his simple constitutive rule and his theory of collective acceptance – only because it is familiar to many readers. I should have been clearer that I meant it merely as an example for illustration, not because I endorsed either of these. In the second half of *The Ant Trap*, I explore how various social facts are grounded, and in doing so highlight the severe inadequacy of Searle's constitutive rules. I also give examples of anchoring that do not involve collective acceptance (for instance, in Chapter 7). It is beyond the scope of this book to develop a full theory of how anchoring works (see Epstein 2014, 2016), but collective acceptance is certainly not adequate.

Some of Hawley's and Mikkola's objections are related to one another. Both object to the key distinction in my model – the distinction between grounding and anchoring. Anchoring is superfluous, they argue, and the framework could better be formulated in terms of grounding alone. Hawley also raises objections to how I formulate the framework for social ontology – in particular, features of what I call 'grounding facts' and 'frame principles'. She suggests that there are problematic ambiguities in some of the terms I employ, and argues against my inclusion of grounding in the formulation of frame principles (an objection she applies to my 'regress' argument). Mikkola further worries about the formulation, focusing her attention on the idea of 'grounding conditions', and raising objections to the use of grounding in social ontology altogether. She also worries about the use of the 'constitution' relation in the book. I will start by clarifying grounding facts, grounding conditions, and frame principles, and then return to the grounding/anchoring distinction and to constitution.

Grounding facts, grounding conditions, and frame principles

Mikkola and Hawley raise worries about a number of components of the grounding–anchoring model, including the notions of 'grounding facts', 'grounding conditions', and 'frame principles'.

Hawley points out that the term 'grounding facts' is ambiguous. I must confess that I did not notice the ambiguity when I was writing the book, and regret the confusion this may cause. I do want to emphasize, however, that it does not generate any problems in the arguments of the book, since I *always* use it in the first of Hawley's two readings, never in the second.

The first reading – the one I consistently use – is this: the grounding facts are the facts that ground, as contrasted with the facts that are thereby grounded. On p. 72, I define the term explicitly – 'the **grounding facts $g_1, \ldots, g_n$** are the facts that ground a social fact f', and I illustrate it in figure 6B. In that figure, an example of a 'grounding fact' is the fact *Billy was printed by the Bureau of Printing and Engraving*. An example of the 'social fact' that is thereby grounded is *Billy is a dollar bill*. This is the consistent way that, without exception, I use the term 'grounding facts'. That is, to mean 'the facts that ground'.

There is, as Hawley points out, another completely different notion that that term could be used to label. Namely, facts about what grounds what, e.g. facts of the form *G grounds f*. And Hawley is correct that some people use the term to mean this different notion (e.g. deRosset 2013). Nonetheless, in the book, I never use the term 'grounding facts' in this way.

The related term 'grounding conditions' raises more trouble, especially in Mikkola's comments. I should have been more explicit that the idea was to use the term in a similar way to more familiar usages of ' … -conditions'. In speaking of the grounding conditions of facts, I meant it to be analogous to other sorts of … -conditions, such as the *truth conditions* of propositions, the *satisfaction conditions* of formulas, the *application conditions* of predicates, the *instantiation conditions* of properties, and the *identity conditions* of sorts or kinds.

In all these cases, the … -conditions are general conditions for an entity to have a certain characteristic (involving whatever fills in the ellipsis). Truth conditions of some proposition P, for instance, are conditions on the world such that if the world meets those conditions, P is true. Application conditions for predicate *P* are general conditions such that if the object designated by a term *d* meets those conditions, *P* applies to *d* (i.e. the sentence '*d* is *P*' is true). Identity conditions for kind *K* are conditions on a pair of objects such that if the objects are both *K*'s and the pair meets those conditions, the first element of the pair is identical to the second.

Grounding conditions should be understood similarly. The grounding conditions of fact *f* are conditions on a plurality of facts *G* such that if *G* meets those conditions, *G* grounds *f*.

Truth conditions and application conditions link propositions, sentences, and predicates with the world. Instantiation conditions link properties with the world, and identity conditions link facts about the identity of objects with other relational facts about objects. And grounding conditions link facts with other facts that ground them (i.e. they link 'grounded facts' with 'grounding facts').

One thing that may be confusing is that I frequently speak of the grounding conditions of a fact-schema, rather than of a fact. For instance, I speak of the grounding conditions of facts like *x is a war criminal* more than I do of facts like *Genghis Khan is a war criminal*. But that is actually common in speaking of ... -conditions. For instance, in semantics we might give truth conditions for the proposition *The ball is red* or for propositions of the form *x is red* or for propositions of the form *x is Y*. For any of these – whether it is a particular proposition or a proposition-schema – there are many ways the world could be that would make the proposition true, and many ways of specifying these conditions.

Similarly for the grounding of facts and fact-schemas. Consider the following fact:

f: This computer screen is now displaying a red square on a blue background.

There is not a unique plurality of grounds for this fact. There is a plurality of facts *G1* about the arrangement of pixels on the screen that grounds *f*; there is a plurality *G2* of facts about the arrangement of various clusters of pixels on the screen that grounds *f*; there is a plurality *G3* about facts about power going through transistors, diodes, etc., that grounds *f*; and innumerably more. There are also other ways *f* might have been grounded but was not actually grounded. We might give conditions for the grounding of *f* in terms of arrangements of pixels, or in terms of clusters of pixels, or in terms of power going through various circuits, etc.

It is not much different to give the grounding conditions for a fact-schema. Instead of fact *f*, we might consider *f'*:

f': *x* at *t* displays a red square on a blue background.

The conditions on a plurality of facts *G* that grounds *f'* will involve *x* and *t*, and will be more general than conditions involving the particular kind of computer screen in *f*. But apart from being more general, it is not much

different to provide a set of grounding conditions for a fact-schema than it is to provide one for a particular fact.

The grounding conditions of a fact are not an analysis of the fact. Rather, a set of grounding conditions for f gives conditions on G such that if G satisfies those conditions, G grounds f. Again, this is similar to other ... -conditions. Consider, for instance, the identity conditions R for a kind K. The statement of a criterion of identity often takes the following form: $\forall x, y((Kx \wedge Ky) \to (x = y \leftrightarrow xRy))$. The criterial relation R is not an analysis of K. It just is the minimal sufficient conditions to guarantee that if x and y are both K's and stand in that relation, then x is identical to y. Likewise, if G satisfies a set of grounding conditions for f, that suffices to guarantee that G grounds f.

The same example helps clarify what a frame principle is. In discussing the criterion of identity, we have discussed two different but related things:

- the criterion of identity: $\forall x, y((Kx \wedge Ky) \to (x = y \leftrightarrow xRy))$;
- the criterial relation: R.

These two are different, but they are of course closely related. Since the criterion of identity takes a canonical form, it is fair to say that to ask about the criterion of identity is precisely to ask about the criterial relation. Still, they are not the same thing.

Likewise with grounding conditions versus frame principles. A frame principle for fact f is like the entire statement of the criterion of identity: it is an explicit statement that expresses the grounding conditions for f. For the fact *x is a war criminal*, for instance:

- the frame principle: *necessarily, if x commits such-and-such acts in the course of a war, then that fact grounds the fact that x is a war criminal*;
- the grounding conditions: *committing such-and-such acts in the course of a war*.

The frame principle is the overall expression of the link between the grounding facts (i.e. the facts that satisfy the grounding conditions) and the grounded fact.

A simpler toolkit, then, might make do with just one of the two – the grounding conditions or the frame principles. But as with criterial relations and criteria of identity, it is sometimes more convenient to speak of one rather than the other. Especially when frame principles get more

complicated, it can be helpful to speak only of the grounding conditions rather than re-expressing the frame principle each time.

I hope this will help clarify certain statements that Mikkola worries about. For instance, it helps explain what I mean when I write, 'To ask about the grounding conditions for facts of the form *x constitutes a mob* is precisely to ask about the frame principles for facts of that form'. That is not to say that the grounding conditions are the frame principle, any more than the criterial relation is the criterion of identity. Rather, it is that providing one is tantamount to providing the other.

It also explains why it cannot be correct, as Mikkola suggests, to characterize me as claiming that 'the frame principle stays the same, but the grounding conditions differ for different worlds'. In my framework, that does not make sense. The grounding conditions for a fact *f* do not change from world to world. The grounds may change from world to world, but the conditions do not. And those (fixed) conditions are expressed in the frame principle.

The distinction between grounding and anchoring

As Hawley points out, making the intuitive discrimination between grounds and anchors helps untangle longstanding problems in social ontology. For instance, it helps distinguish two different versions of individualism about the social world, and helps clarify what's going on in situations where people play multiple roles in social situations they are involved in. I agree with Hawley that the distinction helps clarify ethical features of groups that are 'self-creating' and those of groups whose features are anchored by outsiders. I discuss this sort of case in the second half of *The Ant Trap*; see especially Chapter 17 on Tuomela and Gilbert.

Hawley thus agrees that there is an intuitive and illuminating contrast to be made: we should distinguish the facts that 'set up' or 'construct' social categories (i.e. the anchors) from the 'building blocks' or 'constituents' of social facts (i.e. the grounds). The point on which we disagree is the nature of the anchoring relation: whether it should or should not be regarded as a sub-species of grounding. Hawley and Mikkola both object to my claim that anchoring is best understood as a separate metaphysical determination relation from grounding.

The very idea of introducing anchoring – yet another metaphysical determination relation – may strike some readers as refreshingly bold, but others may find it pointlessly hubristic. Grounding is getting a lot of attention nowadays, but even so, to some philosophers (e.g. Wilson

2014) the grounding relation already seems objectionable. Only very reluctantly did I come to acknowledge the need for a sharp distinction between anchoring and grounding. However, it is the best choice if we are to make sense of the grounding of social facts – and in particular, if we are to get the *modal signature* of social facts correct. That is, if we are to properly capture the possible worlds in which various social facts obtain and the possible worlds in which they do not. In a metaphysics that collapses anchors into grounds, the modal signature of some social facts cannot be properly accommodated.

A good way to show this is by answering Hawley's question about the analogy between modal and temporal cases, with regard to a fact such as *Genghis Khan was a war criminal*. Hawley and I agree on almost everything regarding the modal signature of such facts. But, I will argue, Hawley misses the crucial implication that it forces the distinction between grounds and anchors.

The category *war criminal*, like many legal and social categories, is fairly complex. What it takes for a person to be a war criminal is anchored in part by the enactment of the Geneva Convention and various judicial statutes, but that is not all. The contours of war crimes are also affected by the outcomes of actual trials, by actual practices in wars, by judicial interpretations, and more. Let us symbolize these facts (the facts that anchor the grounding conditions for *x is a war criminal*) as $A_1, \ldots, A_n$, and the plurality of these as A.

The plurality of facts A, of course, obtains in the actual world. And many facts of the form *x is a war criminal* also obtain in actual world. Among these, many obtain in the same geographic and temporal contexts that A does. The war crimes of Bashar al-Assad, for instance, took place in Syria, a signatory to the Geneva Convention. Nonetheless, the category *war criminal* is not limited in its application only to locations in which it and other international treaties have been enacted. Even the citizens of nations that are not signatories to the Geneva Convention or other international treaties can be war criminals. Some social properties obtain only in local contexts, but *war criminal* is insensitive to geographic context.

Likewise for temporal contexts. Though he died in the thirteenth century – well before the Geneva Convention – Genghis Khan was a war criminal, having ordered the slaughter of over a million people in Nishapur alone. This because the conditions for being a war criminal are anchored to apply retroactively as well as prospectively. Though the facts in A are largely facts about twentieth-century Europe, that does not mean that the category anchored by A only applies to those localities and times.

And likewise for modal contexts. Consider the possible world *w* that is just like the actual world except that it ends in the year 1500. In that world, Genghis Khan also ordered the slaughter of Nishapur, and so he was a war criminal, even though there is no twentieth-century Europe in that world at all. Or, to make the point even sharper, consider the possible world *w'* that ends in the year 1500, but where there is another war criminal who does not even exist in the actual world. In addition to the war criminal Genghis, there is also his brother Chuck, who committed similar atrocities. Though the plurality of facts **A** does not obtain in either *w* or *w'*, that does not prevent Genghis and Chuck from being war criminals in *w* and *w'* respectively. In this sense, there is no disanalogy at all between the geographic, temporal, and modal cases.

On this much, Hawley and I agree. We need to make sense of facts in other geographies, times, and worlds that involve the socially constructed categories of those other contexts; and also of facts in other geographies, times, and worlds that involve the socially constructed categories of our own context. (This bears some relation to diagonal versus horizontal propositions in semantics, and also as Hawley mentions, to the 'emic/etic' distinction in anthropology.) Hawley says that I overlook the former (diagonal, emic, internal, or inner) contexts and only discuss the latter (horizontal, etic, external, or outer) contexts. But that is not correct: I discuss both in Chapter 9, and show how the grounding–anchoring framework makes sense of them.

It is true that I focus more on the latter contexts – the horizontal, etic, external, or outer ones. The reason is that these contexts demonstrate the inadequacy of conjunctivism. ('Conjunctivism' is the term I introduce for the view that the anchors are among the grounds, see Chapter 9.) My observation and argument is that conjunctivism can *only* make sense of vertical/emic/internal/inner contexts. If these were the only contexts, there would be nothing to decide between conjunctivism and the grounding–anchoring framework. It is the horizontal/etic/external/outer contexts that are fatal to conjunctivism, and hence that force us to separate anchoring from grounding.

Hawley is of course right that the grounds of some social facts are more sensitive to context than the *war criminal* cases. Many social facts do have geographically, temporally and even modally specific grounding conditions. For instance, the fact *x is an insider trader* is temporally sensitive: it was anchored in 1934 to be prospective, not retrospective. A person who traded on insider information in 1933 is not an insider trader, while a person who traded on insider information in 1935 is. But *war criminal*

is a simpler case. The anchors 'set up' the category for all contexts, and in none of those contexts is it required that the anchors obtain for *Genghis is a war criminal* or for *Chuck is a war criminal* to obtain.

How does this demonstrate that anchoring is *not* a kind of grounding? The facts *Genghis is a war criminal* and *Chuck is a war criminal* are examples of facts whose obtaining does not require that the anchors (or a substitute set of anchoring facts) themselves obtain in that world at all. The modal signature behaves no differently than does the temporal or geographic signature. But if that is right, then anchoring is distinct from grounding, because the modal case has an *implication for grounding* that the geographic and temporal cases do not. Because even though grounding is not geographically bound or temporally bound, it *is* world-bound.

In saying 'grounding is not temporally-bound', I mean this: consider some fact about a particular time, such as *This mark is a footprint at 5:06pm on January 1, 2018*. That is grounded in part by facts about a foot striking the ground at a different, earlier time. In saying 'grounding is not geographically-bound', I mean this: consider some fact about a particular location, such as *The President of the United States delivered the State of the Union address at the Capitol*. That is grounded in part by facts about the person delivering the address being President, which is grounded in part by facts about elections across the United States.

But the same does not hold for grounding across worlds. Categorical (non-modal) facts are not grounded by other-worldly facts: they are grounded by other categorical facts. Grounding is a relation between a plurality of facts that obtain in a world and a single grounded fact that obtains in that world.

This, at least, is the prevailing understanding of the grounding relation. If a plurality of facts G partially grounds fact f in world w, that implies that both G and f obtain in w. But the war criminal case is an example in which a plurality G is part of the metaphysical explanation for f, and f obtains in w, but G does not. Thus that plurality of facts – despite being part of the metaphysical explanation for f – is not a partial ground of f. Anchoring, in short, is distinct from grounding.

The conjunctivist does have avenues for response. For instance, one could revise the prevailing understanding of grounding, or instead one could insist that lots of facts that appear categorical are actually covertly facts about other worlds. But we should not underestimate the costs associated with these avenues.

A further good reason to distinguish anchoring from grounding is that they have different relata. Grounding is a relation between a plurality of

facts *G* and a fact *f*. But anchoring is not the metaphysical determination of a fact. Rather, it is the metaphysical determination of a way facts of a certain sort are grounded. We might think of anchoring as setting up the grounding conditions for a fact or fact-schema. Or similarly, we might think of anchoring as setting up the frame principle that expresses those grounding conditions. Or even more generally, we might think of the complete anchoring relation as standing between a set of facts and the grounding relation as a while. All of these are variants on how we can model the anchoring relation. But whichever of these we favor, anchoring does different work than grounding does.

Constitution

Among the frame principles I discuss in the book are ones involving constitution. The fact *The Supreme Court is constituted by {Alito, Breyer, Ginsburg, Gorsuch, Kagan, Kennedy, Roberts, Sotomayor, Thomas} on January 1, 2018*, for instance, is a very different fact from *The Supreme Court exists on January 1, 2018.* In the book, I argue that even this simple point is overlooked in much of the social ontology literature. For instance, Searle's 'X counts as Y in C' formula and Hindriks's 'status rule' and 'base rule' formulation both collapse such facts into one another – that is, facts about the existence of a social entity and facts about what constitutes a social entity.

Focusing on facts about constitution in particular, Mikkola worries about how we are to understand these. She correctly points out that the term 'constitution' is used in many ways in the philosophical literature, and that this is exacerbated by Searle's term 'constitutive rule'. In the book, however, I try to be clear that I am speaking of material constitution – a relation between two objects that holds when the first is the 'material' of the second. I use it in the sense of the literature on the 'statue and the clay', such as the 'coincidentalist' view that the statue and the clay are distinct objects from one another, even though they occupy the same space at the same time.

In Chapter 10, I discuss constitution at length, and actually propose an explicit analysis of the constitution relation in terms of grounding. For object A to constitute object B at *t*, I propose, is for two conditions to hold: A coincides with B at *t*, and certain facts about A partially ground the fact that B has the material constitution it does at *t*. (For more detail, see *The Ant Trap*, 146–149.) If my analysis is correct, then constitution does not add 'another metaphysical layer' to my ontology.

But even if not, I regard constitution to be a remarkably useful meta-physical tool, especially for making sense of social entities. For instance, it is common for distinct social groups to have the same members as one another, even through the entire duration of their existence (see, among many others, Gilbert 1989; Epstein 2017, and *The Ant Trap*, 139ff.) Even an unanalyzed constitution relation helps untangle this. My aim in the book is not to develop an austere or 'desert landscape' ontol-ogy, even while I do attempt to explain the grounds or foundations for social entities. I aim to draw on the most promising tools in metaphysics to do so. Toolkits change, and the most promising tools of today may not be the ones we will prefer to use tomorrow. But I do not shy away from seizing on the best ones we currently have, if they help distinguish things that are otherwise conflated.

Group agency

Frank Hindriks's interesting comments focus largely on the last two chapters of the book, in which I discuss the 'social integrate' and 'corpor-ate agent' theories of group agency. In those chapters I criticize both models, and as Hindriks points out, I do so for distinct reasons. Among the representatives of the social integrate model are Bratman (2014), List and Pettit (2011), and Gilbert (1989). Key representatives of the cor-porate agent model are Searle and to some extent Tuomela. Hindriks himself endorses both models – some social groups are social integrates, he claims, while others (perhaps a subset of the former) are corporate agents.

I will start by briefly summarizing the idea of a social integrate theory and its flaws. Social integrate theories are largely concerned with making sense of group attitudes and group agency – what is it for a social group to have a shared intention or belief, to be an agent, or to be subject to certain group norms? The theories in this family endorse the following: groups have these properties (having a shared attitude, being an agent, etc.) in virtue of the attitudes and commitments of members of the group toward one another and toward the group as a whole. That is, the members of the group have to be 'integrated' with one another in some way.

Suppose, for example, that group G has intention J. To explain this, a social integrate theory will look to the beliefs, knowledge, intention, and other attitudes of members of G. Michael Bratman's theory is perhaps the clearest example of a social integrate theory. In his view, the fact

that group G has intention J is a matter of various interlocking intentions and common knowledge of members of G.

Many other prominent theories of group agency – including List and Pettit's and one part of Hindriks's – are social integrate theories as well. Hindriks, for instance, argues that part of what grounds the agency of a group is the acceptance by members of a group of a collective decision procedure. This acceptance is a matter of members having certain mental states. Thus Hindriks's proposal amounts to adding certain 'integration' conditions on the members of the group – that they have the appropriate acceptance attitudes – in order for the group to be a group agent.

Social integrate theories of attitudes, agency, and norms are a very general class of theories: it is hard to see what could possibly be the problem with them. How could group agency depend on something *other* than some sort of mental integration – i.e. the attitudes and commitments and other mental states – of group members?

Explaining this is one of the key aims of the second half of *The Ant Trap.* It shows how simple facts about groups are grounded by facts *other* than facts about the members, and other than facts about member attitudes in particular. It gives wide range of examples of facts about groups – what makes them come to exist, what makes them have the membership they do, what makes them have the abilities and powers they do – that depend on external facts. These examples show that mental state integration – no matter how detailed or nuanced – is neither necessary nor sufficient for group action, and neither necessary nor sufficient for group intention. Mental state integration is not necessary: with the right external facts, a group with very little internal integration, and whose members have very few attitudes, can nonetheless have group intentions and take group actions. And mental state integration is not sufficient: it does not matter how many attitudes, acceptances, and other mental states you pack into group members. None of that guarantees that a group has an intention or takes action, if the external contributors are infelicitous. Facts entirely separate from facts about group members can make or break group action and intention.

How is this possible? Consider the deliberation, intention, planning, and action of the group of shareholders of the Microsoft Corporation. A vote taken by the group of Microsoft shareholders, for instance, does involve the attitudes of shareholders, but it also depends on more. Among other things, the group's vote depends on the ownership stakes of each shareholder, and those stakes in turn depend on much more than the member attitudes. Of course, this is just one illustration – this and other

examples are described and argued in more detail in Chapters 15 and 16 – but even with this we can see that the attitudes and accepted decision procedures do not exhaustively determine the result of a vote.

Of course, a group's accepted collective decision procedure might include the provision that votes are weighted by shareholdings. But this is actually to concede the point. Suppose members of a group agree on a procedure in which the outcome of the vote does not supervene on their own attitudes. That implies that the outcome of the vote does not supervene on member attitudes. Merely because their agreement involves their attitudes does not mean that the result of the vote is determined by those attitudes.

It is external factors like these that prevent mental state integration from being either necessary or sufficient for group intention and action. We often set up groups to accommodate the fact that member attitudes and other mental states are unlikely to be integrated, with external factors filling in the gap. Conversely, where attitudes would seem to be sufficient, external factors can frustrate or defeat a group having an intention or taking an action.

What does it take, then, for a group to be a group agent? In *The Ant Trap*, I actually begin with the same underpinnings of group agency that Bratman does, and that List and Pettit agree with. Namely, group agency should be understood functionally. We understand group agency to be functionally analogous to individual agency. A widely accepted theory of individual agency is that people have a modular system that we use to navigate the world. We have beliefs, knowledge, form intentions, plan, reason, and take action in the world. All of these components perform their own functions, and they integrate with one another to perform unified functions. They compose, in other words, a 'system of practical activity'. Group agency is understood on this model: we look for roughly the same functional system but as realized by groups rather than by individuals. I follow this approach to group agency in *The Ant Trap*. The problem is not that functional underpinning; rather, it is that its proponents have not been consistent in applying it.

Hindriks agrees with me that Bratman does not even attempt to provide an account of group agency. But I go further. Even as a theory of shared intention Bratman's account does not work. Part of the reason is what I just discussed: the account incorrectly assumes that mental state integration is sufficient for a group to have a shared intention. Another equally significant reason, however, is that shared intention cannot be analyzed in isolation from the analysis of group agency as a package.

On a functionalist understanding of practical activity, it is an illicit move just to pull out one component of that system and try to work out what will 'realize' that component. That is not how functional realizations work. In looking for the realizer of a functional system, you need to realize the entire system. How groups realize the 'intention component' of the system depends not just on internal features of that component, but also on how that component interacts with the other components. And that interaction depends on the actual realizers of those modules. In theories of group agency, we look for kinds that realize group intention, group action, group belief, group deliberation, and so on, all together as realizers of a group system of practical activity. Any given set of realizers needs to be assessed as realizers of the whole system.

Hindriks, Bratman, and others seem to want to sever the analysis of group intention, or group belief, or group deliberation, from group agency. But how do we make sense of this? What conditions do group intentions need to satisfy, if not to play an appropriate role in a group system of practical activity? We either need to apply the functionalist treatment consistently, or else we need an alternative understanding of group attitudes and action altogether.

Being clear on this point helps explain my challenge to the 'status account' of 'corporate agency'. Hindriks correctly interprets me as criticizing the idea that groups are assigned the 'status' of being agents. One of Hindriks's main objections to me is that it is incorrect to regard status theorists as assigning agency. Hindriks argues that there are many important statuses that do get assigned to groups of various kinds, and that I pay insufficient regard to these.

I agree with Hindriks that groups of many kinds have powers, capacities, and norms. I also agree that these are pertinent to group agency. This, however, does not motivate the status account. At least three features of the status account are problematic: that statuses are 'assigned' by collective agreement or acceptance; that powers, capacities, and norms are best understood as 'statuses'; and the question of whether all group agents – including corporate agents – have to be social integrates.

To begin with the last of these, Hindriks surprises me with his endorsement of the view that all agency is a matter of social integration, and that for a corporate agent to be an agent requires that it be a social integrate. He points out that Searle does seem to regard the assignment of agency as a kind of status assignment, and he departs from Searle on this point. I sympathize with the distancing from Searle, but it had been my impression that one of the key attractions of a 'corporate agent' picture

was to explain the nature of typical groups like corporate and university boards, teams, committees, and so on. These often fail to meet the conditions of even an accommodating 'social integrate' theory, in that members may have clashing attitudes and not even be aware that there are externally imposed decision procedures, and yet still seem able to be group agents. That is, I had thought that corporate agency served – at least in part – to explain how some burdens of mental state integration could be offloaded from group members, because of powers, capacities, and norms assigned to the group by others.

But this apparently is not Hindriks's view. In his comments, he argues that corporate agents are a subset of the social integrates. If this is right, then my criticisms of mental state integration as giving sufficient or necessary conditions for group agency apply as much to Hindriks as they do to Bratman and List-Pettit.

In contrast, I do think that the burdens of mental state integration are often offloaded from group members. The Supreme Court, for instance, is a group agent even in the absence of much mental state integration. But the way that is done is not with assignments of status. Instead, as I discuss in Chapters 11–13, the Supreme Court is anchored to have certain powers, capacities, and norms. Those are not statuses, and they are not anchored by collective acceptance. In *The Ant Trap* I do not develop a full theory of anchoring, but I do discuss some heterogeneous ways that anchoring takes place. For instance, I examine how laws are anchored (see Chapter 7), and discuss the many factors that figure into their anchoring.

Powers, capacities, and norms play an important role for the Supreme Court. In fact, I regard them as playing an even more central role than Hindriks does. Facts like when the Court is in session, and what the Court has jurisdiction over, are not just appurtenances to the Court as an agent; instead, they are part of how the Court realizes a system of practical activity. They are part of what takes up the burden of realizing the system even when the member attitudes fail to suffice. And these facts are anchored to be grounded by external features of the world. If you only looked in the heads of the Justices, they would not be sufficiently integrated to be a group agent. But if you look at the wider set of facts that ground these properties, you can see that it realizes a system of practical activity.

The Supreme Court can be a group agent without being a social integrate. This is in part because of the powers, capacities, and norms that it is anchored to have – they are part of why the group satisfactorily realizes a system of practical activity. The Court is not an agent-qua-social-

integrate that then has powers, capacities, and norms layered on top of it. Moreover, these powers, capacities, and norms are not status assignments, nor are they anchored via collective acceptance.

Just how loosely affiliated can a group of people be, in order to be a group agent? In the case of a class action lawsuit, it is unlikely that sufficient structure is anchored for the class to be a group agent. (In fact, this is reflected in American law on class actions, where the action is regarded as being taken by a law firm on behalf of a class, rather than being an action of the class.) The class can, of course, be a 'legal entity', but that is not the same as an entity that takes an action. In the general case, the necessary degree of integration depends entirely on how much external structure the group is anchored to have. To be a group agent – at least under the functionalist understanding of group agency – a system of practical activity needs to be realized. For some groups, the bulk of that realization is done by mental states of group members; for others, the bulk is done by external facts. (For more on the nature of groups and their norms, see Epstein [2017].)

Let me conclude with a clarification about memberless groups. I deny, rather than endorse, the claim that groups cease being objects when they have no members. To use the example of the Supreme Court again, it can persist even if all the members are dismissed and subsequently replaced. I explain in Chapter 12 how we can identify groups even during the times they are memberless. Here too, it is not 'status assignments' that explains this occurrence, but rather (as I discuss in the surrounding chapters) how their existence and constitution conditions are anchored.

Disclosure statement

References

Bratman, Michael. 2014. *Shared Agency*. Oxford: Oxford University Press.
deRosset, Louis. 2013. "Grounding Explanations." *Philosophers' Imprint* 13 (7): 1–26.
Epstein, Brian. 2014. "How Many Kinds of Glue Hold the Social World Together?" In *Perspectives on Social Ontology and Social Cognition*, edited by Galloti Mattia and Michael John, 41–55. Dordrecht: Springer.
Epstein, Brian. 2015. *The Ant Trap: Rebuilding the Foundations of the Social Sciences*. New York: Oxford University Press.
Epstein, Brian. 2016. "A Framework for Social Ontology." *Philosophy of the Social Sciences* 46 (2): 146–176.

Epstein, Brian. 2017. "What Are Social Groups? Their Metaphysics and How to Classify Them." *Synthese*.

Gilbert, Margaret. 1989. *On Social Facts*. Princeton, NJ: Princeton University Press.

List, Christian, and Philip Pettit. 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.

Wilson, Jessica. 2014. "No Work for a Theory of Grounding." *Inquiry* 57 (5–6): 535–579.