



A University of Sussex DPhil thesis

Available online via Sussex Research Online:

<http://sro.sussex.ac.uk/>

This thesis is protected by copyright which belongs to the author.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Please visit Sussex Research Online for more information and further details

**Plural Bodily Subjects:
A Radical Account of Thinking and
Acting Together**

Richard Andrew Weir

PhD Social and Political Thought

University of Sussex

November 2015

Declaration

I, Richard Weir, hereby declare that this thesis has not been, and will not be, submitted in whole or in part to another University for the award of any other degree.

Signature..... Date.....

University of Sussex

Richard Andrew Weir

PhD Social and Political Thought

Plural Bodily Subjects:

A Radical Account of Thinking and Acting Together

Summary

The primary aim of this thesis is to defend the idea that there are ontologically collective forms of thought and action. This is to say, that there are at least some instances in which a thought or action is appropriately ascribed not to the individual members of a group, but only to the group as a whole. In chapters 2 and 4 existing attempts, primarily in analytic philosophy, to defend such phenomena by appealing to either the content, mode, or subject of intentional states are criticised. These criticisms in turn motivate an alternative understanding of subjectivity, outlined in chapters 3 and 5. This alternative draws on the phenomenological work of Dan Zahavi and Maurice Merleau-Ponty, to argue that subjectivity must be understood, firstly, as constituted by the pre-reflective self-awareness that is central to all intentional experience and, secondly, as intrinsically bodily. Finally, in chapter 6, and by drawing on Merleau-Ponty's thoughts on habituation and intersubjectivity, it is argued that it is possible to understand groups as continually in the process of developing such a form of plural bodily subjectivity through processes of group-level habituation. Overall, therefore, a radical position will be defended, which holds that not only can groups think and act in an ontologically collective sense, but that they can do so in virtue of the fact that they can achieve a certain level of phenomenal self-consciousness. However, this position will be tempered by the thought that unified self-awareness and subjectivity is a matter of degree; where to have a unified pre-reflective sense of self is to be an individual subject, groups must be understood as always in the process of developing a form of unity that is, ultimately, out of their grasp.

Contents

Acknowledgements	6
Chapter 1: Introduction.....	8
I: Linguistic and Ontological Collectivity.....	12
II: The Priority of the Social.....	14
III: Contemporary Anti-Individualism.....	17
IV: Phenomenology and the Philosophy of Mind	21
V: Overview.....	22
Chapter 2: The Intentionalist Approach	27
I: Collective Intentional Content	29
II: Collective Intentional Mode	32
III: The Ambiguity of the First Person Plural	38
IV: The Problem with the Intentionalist Approach	43
V: Conclusions	47
Chapter 3: The Subject and Self-Awareness.....	49
I: Shoemaker and Non-Perceptual Self-Awareness	50
II: Kant and Pure Apperception	54
III: Zahavi and Pre-Reflective Self-Awareness	61
IV: The Unity of the Self.....	65
V: Conclusions	70

Chapter 4: Four Accounts of Plural Subjecthood	72
I: Bacharach.....	73
II: Rovane	81
III: Gilbert	91
IV: Schmid	98
V: Conclusions	106
Chapter 5: Bodily Self-Awareness.....	107
I: The Subject-as-Embodied View.....	110
II: The Subject-as-Bodily View	117
III: The Spectrum View	124
IV: Conclusions	137
Chapter 6: The Plural Bodily Subject.....	140
I: Merleau-Ponty on Habit and Extension	142
II: Merleau-Ponty on Intersubjectivity	146
III: The Problem of Self and Other	150
IV: Plural Objective Self-Awareness.....	152
V: The Group and Habituation	157
VI: Objections and Responses	162
VII: Final Conclusions	169
Bibliography	172

Acknowledgements

The success of any thesis is never that of the individual. If anything of any interest is found in the following pages, it is only because I have had the good fortune to be surrounded by the right people to give me a shove in the right direction when I needed it.

To begin with, I have enjoyed the support of an incredible community of postgraduate students over the last few years. They have undoubtedly had a profound effect on my work. More importantly, however, they have been the source of many a bizarre and absurd conversation at those moments when the last thing I wanted to think about was philosophy. In particular, my thanks go to Jacob Berkson, Tim Carter, Ezra Cohen, James Cullis, Robb Dunphy, Adrian Downey, Alex Elliott, Jana Elsen, Chris Ferguson, Alastair Gray, Christos Hadjioannou, Phil Homburg, Dimitri Kladiskakis, Valentinos Kontoyiannis, Patrick Levy, Gabriel Martin, Chris O’Kane, Huw Rees, Elliot Rose, David Martinez Rojas, Alan Roberts, Umut Sahverdi, Arthur Willemse, and Joel Winder.

My thanks also go to the faculty of the Department of Philosophy and the Centre for Social and Political Thought at the University of Sussex. I have been at Sussex, in one form or another, for almost a decade in total now, and there is not a single member of faculty who has not taught me something in that time, whether that was in a classroom or in the pub.

In particular, I cannot begin to express the debt, both academic and personal, that I owe to my supervisor Andrew Chitty, who has guided me in one way or another since the moment I stepped foot on Sussex campus. If a single word of the following makes sense or is seen to form part of a cogent argument, all the credit belongs to Andrew. Conversely, any faults are my own and remain there in spite of his tireless perseverance in attempting to bring order to the chaos.

Finally, my thanks go to my family. In particular, to my parents, Gwynneth and Keith, for their constant support and for being gracious enough to proofread their son’s terrible writing. To Jess, for her loving companionship. And, more than anything, to my wife Shraddha, for her unwavering ability to keep me afloat. If there was any justice in the world, her name would be on the front of this thesis alongside my own. However, in this instance, I am glad that this would not be appropriate, because I would hate to give her the satisfaction!

To the memory of
W. Gordon Wright

Chapter 1

Introduction

From the perspective of common sense and intuition it would seem that there is nothing particularly problematic about those occasions in which we think or act as part of a group. As social creatures we think and act together on such a regular basis that we rarely pay much attention to our doing so. And yet, it is perhaps precisely this fact that makes it so interesting from a theoretical standpoint. We are presented with phenomena that are so intrinsic to the leading of our everyday lives, so basic to what it is to be human, that we rarely consider what is actually going on under the surface. Nevertheless, once we do begin to scratch this surface, even a little, we find that this quite fundamental feature of human social reality is, in fact, hugely problematic. Like so many things that we take for granted in everyday life, under the glare of sustained analysis we quickly realise that there is virtually nothing that we can take for granted with respect to thinking and acting together¹.

While it is uncontroversial that we think and act together in at least some sense, whatever this might entail, the purpose of this thesis is to enquire as to whether we ever do so in a quite radical sense, such that its very possibility is doubted by many. Specifically, it is concerned with that sense in which the individuals involved might be said to be part of a larger whole in which their individuality is, at least to some extent, subsumed. To put it differently, the question that this thesis will attempt to answer is whether or not groups are capable of thinking and acting in something like the same sense that we assume individuals are.

Perhaps even the idea that individuals might, from time to time, shrug off the shackles of their individuality and merge with one another, in some sense, is not particularly controversial from the perspective of intuition. After all, even the *Spice Girls* sang about that magical moment, in the midst of love and passion, ‘when two become one’. Of course, they were probably only speaking metaphorically, but it does seem significant that the idea that there might be forms of agency larger than the individual has become an accepted part of the public psyche. Indeed, it would be fair to say that, even as a metaphor, it is one of that recurs often and consistently in the history of human

¹Indeed, the placing of these two phenomena side-by-side assumes that there is a tight link between the two, but even this link is the subject of an extensive literature and is open to sceptical challenge. However, for the sake of space, the intricacies of this link will not be dealt with at all below. Rather, it will merely be assumed that there is some link (in the sense that only creatures that think can *act*, while an object devoid of thought can merely *move*) whatever it might consist in.

thought. For, as every philosopher knows, the Spice Girls' musings on the matter were prefigured by visions such as Plato's *Republic* or St Paul's notion that we might all join as One in the body of Christ.

Nevertheless, what might be relatively uncontroversial from the perspective of the social imagination is far from uncontroversial from the perspective of philosophy and wider academic thought. Indeed, many who are involved in the contemporary discussion surrounding these issues are fond of pointing out that some form of individualism has been the pervasive ontological assumption of both philosophy and the social sciences over the greater part of the last century. While this might be a slight exaggeration, tending as it does to downplay the significance of critical social theory, not to mention many of the insights of phenomenological thought, there is something to the view that individualism is the ontological challenge of the twentieth century. Individualism casts itself as a sort of ontological bedrock. Thought and action, regardless of what one thinks they might consist in or how they manifest in the world, are, it seems, properties of a subject or agent, that is, a conscious being. And individuals, at least on the face of it, are the only entities appropriate to describe as such. From this it follows that only individuals can think and act². Just as an arrow shot through the air neither thinks nor acts but simply moves, so the army shooting the arrows does not, strictly speaking, think or act. For, neither the arrow nor the army are conscious beings. Of course the army, unlike the arrow, is constituted by conscious beings, but this is the very point of individualism: that we ought, therefore, to be able to describe the so-called *thought* or *action* of the army in terms of the distinct thoughts or actions of its constitutive individual soldiers. It is the individuals that think and act, and any such attribution to the army as a whole must be understood as merely a summation, or tallying-up, of their individual thoughts and actions. Thus, for instance, while it might seem that an army expresses its courage or determination in repelling an enemy charge, it actually does neither. There may well be courage and determination on display in such an incident, but it is merely the courage and determination of the individual soldiers involved.

² This is, essentially, the central tenet of *ontological individualism* and it is worth noting that whenever the term 'individualism' is used below, it is intended in this sense. It is important to distinguish this view from *methodological individualism*, which claims that the most detailed or accurate description of any social phenomenon will be one that only makes recourse to the particular thoughts or actions of individuals. However, it will also be assumed that a core feature of methodological individualism is, in fact, a prior commitment to ontological individualism. In other words, while it is possible that one might be an ontological individualist but not a methodological individualist (in the sense that one might believe that holistic or functional forms of social explanation can be useful even if they are not entirely accurate), it is extremely unlikely that one would be a methodological individualist but not an ontological individualist (as this would involve the thought that there are ontologically collective forms of thought or action, but that it is not helpful to appeal to these in social explanation). As such, where the views of methodological individualists are considered below, it is assumed that such theorists are thereby committed to ontological individualism (for a more detailed discussion see Lukes (1978)).

Equally, it is the case that the enemy is repelled, not by the army per se, but by the cumulative efforts of its soldiers.

Similarly, from the standpoint of individualism *the army did not surround the castle*, rather each of the individual soldiers took a certain position in relation to it such that, once situated, the castle was effectively surrounded by soldiers. *NASA did not put a man on the moon*, rather each of the individual members of NASA performed their own specific job such that the end result was a spacecraft, several trained astronauts, and a reasonably good plan so that they might all end up 384,400 kilometres from Earth. *Football teams do not want to win football matches*, rather each individual player wants to win the match. Or perhaps, in some cases, what they really want is personal fame and fortune, but in the pursuit of such individualistic goals manage to give a fairly compelling rendition of a team that wants to win a match. And the list could go on indefinitely, liable as we are to ascribe the full range of human mental capacities and behaviours to groups, no matter how big or small: for the individualist, a certain error is made in every such case. In a pragmatic sense it is often much more sensible to talk in terms of the thoughts and actions of groups, and the individualist can accept this. After all, it would be absurd if every time I wanted to discuss the moon landing, that instead of saying 'NASA did such and such', I had to provide a detailed breakdown of the role of each NASA employee. Nevertheless, what is significant for the individualist is that such a breakdown could be given and that it would give a more accurate description of what actually occurred. To suggest otherwise is to be committed to the idea that groups can be agents or subjects in the same sense as individuals and this, the individualist believes, is absurd.

It might be argued at this point that the mere fact that we regularly make recourse to such group-level descriptions is telling and significant in its own right. Surely there must be something underlying this constant talk of the thoughts and actions of groups? But, the individualist might respond, that we are liable to describe groups in such agential terms should be of little consequence, considering how willing we are to discuss inanimate objects in precisely the same terms. After all, I regularly admonish the stupidity of objects for their petulance in getting in the way of my foot when I have the poor fortune to stub my toe (and if *Disney* movies are anything to go by, this ability to personify the natural world knows no bounds), but I have no desire to attribute actual agency to the corner of my bedstead.

What, then, is wrong with accepting the individualistic perspective? Can one not just say that, while talking in individualistic terms might be cumbersome on an everyday basis, it nevertheless offers a solid basis from which to measure the accuracy of sociological claims? One immediate worry might be just how much social theory we would have to discard if limited to a strictly individualistic

analysis. Would we, for instance, be able to make sense of such influential ideas as Rousseau's *general will*, Hegel's *spirit*, or Marx's *class consciousness* on this basis? Some individualists might be willing to respond that this is so much the worse for such theories; if they are to prove ontologically sound, they must be capable of being comprehended in purely individualistic terms, otherwise they ought to be consigned to the flames. A stance that has, somewhat inevitably, led directly to such attempts being made. A particularly striking example being Jon Elster's attempts in the 1980s to interpret Marxist thought through the lens of methodological individualism (for instance see Elster (1985))³. However, one need not accept this methodological principle in order to remain convinced by the ontological claim that only individuals think and act (see footnote 2). In other words, one can be an individualist and still accept that forms of holistic social explanation will often prove to be the most useful and, therefore, need not be consigned to the flames.

A more compelling reason to reject individualism is the argument that it simply doesn't work. This is to say, that when we do try to apply a strictly individualistic analysis this leads to results that do not match reality or that do not seem capable of capturing certain important social phenomena⁴. A pertinent illustration of this potential shortcoming is provided by rational choice theory, which is perhaps the approach that takes individualism to heart more than any other. Indeed, it is possible to view rational choice theory as committed to only two core premises: (1) that only individuals think and act⁵, and (2) that said individuals act on the basis of means-end rationality. However, rational choice theory has long since been plagued by so-called coordination problems, epitomised by games such as *The Prisoners' Dilemma* and *Hi-Lo*. In broad terms, the problem is that rational choice theory predicts a certain individualistic or non-cooperative pattern of response in such games (or in the case of *Hi-Lo*, no pattern of response at all), but when carried out in experimentation a more cooperative response is witnessed a significant amount of the time. In other words, an approach in which individualism represents a core premise yields results that do not match reality. Of course, there are numerous possible ways that one could respond to such coordination problems. The committed individualist could, for instance, call into question the particular conception of rationality in play or, at least, amend it with conditions that would allow rational choice theory to

³ Although the failings of such attempts, coupled with the lack of support that they have garnered in the long run, is reasonably compelling evidence that this might have been a misguided approach.

⁴ For a fairly brief, yet compelling, argument against individualism along these lines see Searle (1990).

⁵ Phrased in this way, premise (1) expresses a commitment to ontological individualism. To some this might seem strange, as traditionally rational choice theory is understood as committed to methodological individualism. However, methodological individualism should be taken to represent a stronger claim than ontological individualism and, on a common sense view of things, involves at least a prior commitment to ontological individualism (see footnote 2). Moreover, the coordination problems associated with rational choice theory, and which are of primary interest here, appear to arise just in case one accepts (1) and (2).

account for certain cooperative responses. However, a compelling alternative is to suggest that these problems raise doubts over the idea that individualism can provide an adequate account of the nature of all social thought and action⁶.

In truth, however, the debate over individualism has been intense and there have been some compelling defences of individualism against such attacks. As such, it could quite easily occupy another entire thesis in order to provide a truly convincing rejection of individualism on its own terms. Instead, the present thesis can be taken to offer a far more substantial attack on individualism, by way of providing an alternative. After all, the most compelling way in which to demonstrate the inadequacies of individualism is to place the ascription of thought and action to groups on as sound an ontological footing as the ascription of thought and action to individuals.

I: Linguistic and Ontological Collectivity

It is important, however, to make clear exactly what such an anti-individualist alternative is committed to, as there is plenty of room for confusion with such matters. A level of clarification can be achieved by considering a couple of different ways in which the statement of a thought or action involving two or more people might be interpreted. For instance, the statement ‘Tim and Tom are lifting the piano’ appears to be ambiguous between a *distributive* and a *collective* reading in both a *linguistic* and an *ontological* sense. It is very often the case that these ambiguities and, crucially, the differences between them, go unrecognised.

For a start, it must be said that the linguistic ambiguity need not concern us a great deal here. It refers to the fact that in the English language such a statement might refer to two completely different scenarios. The first scenario is one in which Tim and Tom are, perhaps somewhat strangely, taking it in turns to lift the piano. This reading is *linguistically distributive* in that Tim and Tom are performing a single action only in the sense that they happen to be performing a qualitatively identical action, but at different times and not necessarily with any kind of reference to one another. The second, far more likely, scenario is that in which they are lifting the piano together, perhaps by each taking a different end of the piano and thereby sharing the load. This reading is *linguistically collective* in the sense that it refers to a single action event to which Tim and Tom are both party, rather than separate actions that happen to share a type identity. On the distributive reading it follows from the statement ‘Tim and Tom are lifting the piano’ that ‘Tim is lifting the piano’ and ‘Tom is lifting the piano’, but this doesn’t follow from the linguistically

⁶ Many of the issues surrounding rational choice theory and game theory will be discussed in much greater detail in Chapter 4.

collective reading. Thus, with regards to this level of ambiguity it should be reasonably clear that, at the very least, the anti-individualist is only concerned with thoughts or actions that can be considered to be linguistically collective.

It is the second, ontological, ambiguity that is of much greater concern. It refers to the way in which, even if we agree that we are referring to a linguistically collective action, 'Tim and Tom are lifting the piano' is still open to two different analyses of the nature of the action taking place. On the one hand, to take an *ontologically distributive* view in this respect is to suggest that such a statement is in fact a sort of shorthand for what is really going on. This is to say, that when we say that Tim and Tom are lifting a piano, what we cannot mean is that they are the single subject of the action, for they are not a single subject. Rather, they are distinct individuals and, as such, they can only act as individuals. Therefore, a more accurate description of how the piano is lifted would be to say that 'Tim is lifting his side of the piano and Tom is doing likewise'. In the light of the discussion above, individualism can be identified as that position which is committed to the idea that all linguistically collective statements about thoughts or actions describe ontologically distributive states of affairs. In other words, while it might be useful to speak of Tim and Tom's single action of lifting the piano, such talk will always be slightly fictitious and not the entire ontological story, as it were. Similarly, when we talk of *armies* surrounding castles, *space agencies* putting men on the moon, or *football teams* wanting to win matches, these all represent linguistically collective claims about ontologically distributive thoughts or actions; in all of these cases the words that we are employing, in fact, refer to a summation of the totality of individual thoughts or actions that were involved.

On the other hand, the anti-individualist is committed to the idea that, at least on some occasions, such an ontological reduction will not be appropriate or even possible. On this view, when we say that Tim and Tom are performing an action together, we need not just mean that the two individuals are performing two ontologically distinct actions that *add-up* to the lifting of the piano; Tim and Tom's action might not be analysable in terms of its individualistic parts. In this sense the linguistically collective action statement can be understood as referring to an *ontologically collective* action.

Perhaps the clearest way of expressing the difference between the two, is to consider how they would respond to the question, 'who lifted the piano?' The individualist, provided that they were concerned to provide a full answer, would have to say something like, 'well, Tim did his part and so did Tom.' In other words, there is no single answer to the question because, in the case of Tim and Tom, there is no single actor responsible for the lifting the piano. Actions are those of

individuals and there are two individuals involved. The anti-individualist, however, wants to argue that it is sometimes accurate, depending on the particular circumstances, to simply respond, ‘Tim and Tom lifted the piano’, where ‘Tim and Tom’ ought to be read as [Tim and Tom] rather than [Tim] and [Tom]. They did it as a group and the resulting action cannot be attributed to them as individuals, but only to the group as a whole.

There is a sense, therefore, in which to say that an action or thought is ontologically collective is to say that it can be attributed to an ontologically collective subject or agent. This is to say, that what distinguishes the ontologically distributive from the ontologically collective understanding of ‘Tim and Tom are lifting the piano’, is the number or subjects to whom the action is being attributed. If it is anything less than two, then an ontologically collective understanding is being employed. Although this is complicated by the fact that on some of the contemporary anti-individualistic positions considered below, it is argued that there can be forms of action that are only appropriately ascribed to the group as a whole, rather than the constitutive individuals, while it is denied that there is an ontologically collective subject of that action in any literal sense.

Nevertheless, the important point is that when it comes to thinking and acting together, the individualist and the anti-individualist are distinguished from one another by whether or not they accept that some of these thoughts or actions will be ontologically collective. Individualism is the view that all forms of thinking and acting together must be understood as ontologically distributive, that is, as attributable to individual subjects. Conversely, anti-individualism, at least in this context⁷, is the view that there will be at least some occasions when such an understanding is misguided and where the thought or action taking place must be understood as ontologically collective⁸, regardless of the basis on which this is held to be possible and the degree of reality that is afforded to the ontologically collective subject of that thought or action.

II: The Priority of the Social

It might be thought, however, that the above discussion presupposes a certain view of the relation between the individual and society that is suspect. As Frederick Schmitt puts it:

⁷ Someone who accepts that only individuals think and act, yet denies that only individualistic analyses are appropriate forms of social explanation, might reasonably describe themselves as a *methodological anti-individualist*. It should be clear that we are here only concerned with the stronger ontological form of anti-individualism, which denies that only individuals think and act.

⁸ While this thesis will defend the idea that there are ontologically collective forms of thought *and* action, it should be noted that one could be an anti-individualist about actions but not about thoughts, and vice versa. Some of the accounts considered below (most notably Searle’s) appear to accept that there are ontologically collective forms of action while denying that this requires ontologically collective forms of thought.

‘Underlying the debate between individualists and holists is an assumption, questioned long ago by some philosophers influenced by Ludwig Wittgenstein ... that we can understand individual human beings independently of social relations and collectivities. This assumption has been denied on the ground that individual human beings are already bound up in social relations and collectivities merely in virtue of having such attributes as thinking, acting, and speaking a language.’ (2003, 1)

Schmitt's point is that both individualism and the kind of anti-individualism that will be pursued in the course of this thesis, assume that the individual can be understood independently of the social womb in which the individual comes to fruition, but this is an assumption that has been attacked repeatedly over the course of twentieth century philosophy. To put it differently, there is a much more fundamental response to individualism, one which also seems to undercut the need to argue for the possibility of ontologically collective thoughts or actions, which is to argue that the social is prior to the individual (let us call this the *social priority thesis* (SPT)). If SPT is correct then there is good reason to think that individualism is quite straightforwardly mistaken and one does not need a further defence of ontologically collective forms of thought or action to show why.

It should be acknowledged that the idea that sociality is prior to individuality is one which has been fostered well beyond the realms of those influenced by Wittgenstein. For instance, and more pertinently with regards to the direction of this thesis, some form of SPT appears to be almost unanimously accepted within phenomenological thought, exemplified by Heidegger's assertion that there is a sense in which *Mitsein* (being-with), is prior to *Dasein* (being). Furthermore, there can be little doubt that some of those involved in the contemporary debate surrounding anti-individualism completely accept Schmitt's intuition that SPT undercuts the need to defend the possibility of ontologically collective forms of thought and action. For instance, Margaret Gilbert devotes a fairly lengthy chapter of her seminal *On Social Facts* (1989, 58–145) to discrediting a version of SPT, which she attributes primarily to Peter Winch and his interpretation of Wittgenstein's *Philosophical Investigations*, in order to clear the ground for her own anti-individualist argument. In other words, Gilbert's thought appears to be that if the social is in fact prior to the individual, then individualism has been shown to be false at a much more fundamental level and an account such as her own is unnecessary; there is no point in talking about ontologically collective forms of thinking and acting together, if to think or act as an individual is already, in one form or another, dependent on intrinsically social forms of thought and action.

This is, however, a mistake. It is a mistake that is premised on a failure to properly draw a distinction between the ideas of sociality and collectivity, where SPT is concerned with the former and the latter pertains to the forms of thinking and acting together with which, for instance, this thesis is concerned. Of course, it is difficult to provide a definition of sociality that is both accurate and

illuminating. However, in this context, and for the sake of simplicity, it might be understood in terms of the way in which our intersubjective lives are structured in terms of certain rules, roles, and norms. On the other hand, collectivity, as has just been suggested, relates to that sense in which two or more individuals might think or act in concert or unison, that is, together. Of course, these are somewhat artificial definitions (the terms in question have certainly been used in different ways by different authors), but it is fair to say that a distinction along these lines can be made. Now, with this distinction in mind, it should be apparent that the question of whether there are ontologically collective forms of thought and action arises just as long as one accepts that there are such things as individuals. This is to say, that as long as there are individuals, and individual thought and action, then, regardless of the basis on which such individuality is formed, the question arises as to whether those individuals only ever think and act qua individual, or whether they ever join together to think and act in such a way that cannot be understood as ontologically distributive. In other words, SPT is a thesis about the formation of individuals and the question of how individuals are formed is, strictly speaking, tangential to the question of whether those individuals are capable of joining together to think or act in ontologically collective ways. While it may well be the case that individuality only develops out of, and in virtue of, its social context, this does not yet say anything about the capacity of those individuals to think or act collectively.

Thus, what is really at stake with regard to accepting or rejecting SPT is not the necessity or importance of a defence of ontologically collective phenomena, but merely the scope of such a theory. If one takes the phenomenon of thinking and acting together to be concerned with the nature of sociality *per se*, as Gilbert and notably Searle (1995; 2010) seem to, then it would appear to be the case that one does need to first reject the idea that sociality is prior to individuality in order to avoid falling into the trap of an obvious and vicious circularity. For, thinking and acting together in the collective sense must be understood as only arising out of and presupposing a layer of individuality. If sociality is prior to individuality, it would be absurd to then claim that the social realm is defined by some relation between individuals. However, one can still take sociality to be prior to individuality and argue for a level of collectivity that might arise on the basis of that individuality, as long as one is willing to accept that there are forms of sociality other than such collectivity. Consequently, while the present thesis must be taken to be concerned with forms of thought and action that it would not be possible to understand the social world in its entirety without, it should not be taken to be concerned with the nature of sociality *per se* and is simply agnostic when it comes to the accuracy of some form of SPT.

III: Contemporary Anti-Individualism

With these distinctions and clarifications in place, let us turn to those accounts that have attempted to defend the possibility of ontologically collective forms of thought and action. While anti-individualistic theories are by no means novel, in the last 20-30 years there has been a renewed and sustained attack on the assumptions of individualism and numerous attempts made to establish an ontologically sound alternative. This attack has been waged from a range of angles, drawing on the perspectives of academics working in philosophy, the social sciences, cognitive science, and game theory (among others). It is nevertheless united by the attempt to take the possibility of ontologically collective phenomena seriously and subject the ontology of thinking and acting together to a sustained analysis. In many ways, therefore, there is also contained in this attack an implicit reprimand of the past along the following lines: while most political and social theory was willing to talk in terms of group thought and action, the implicit ontological commitments involved in such talk were rarely explicitly defended. While individualism is not the lens through which to re-interpret these positions, there is still a need to be able to cash out their underlying ontological claims.

This wave of contemporary anti-individualism is often referred to, in all its forms, under the heading of *collective intentionality*. However, while useful as an umbrella term that unites a certain field of research, this field has now become so broad that the term has become fairly redundant as a way of designating any particular position, or even set thereof. As such, I will avoid using it for the most part. Instead, I suggest that we can identify three broad, yet distinct, versions of anti-individualism that have arisen in this contemporary debate.

The first approach is difficult to provide with a pithy label, so here it will be referred to it as the *intentionalist approach*, and it is that which is most commonly associated with the work of John Searle, Michael Bratman, and Raimo Tuomela. The reason for this label is that all of the theorists concerned argue, in one way or another, that the key to understanding how and when people think and act together lies in the nature of the individual intentional states of those involved and/or the relations between those states⁹. Of the three approaches in question, therefore, it is perhaps the one that remains closest to the claims of individualism. On this account there are, strictly speaking, only individual subjects and only individual subjects can think in the sense of being conscious and

⁹ When the term ‘collective intentionality’ is not used in the broader sense mentioned above, it is often used in a specific way to designate this approach in particular. However, this is misleading considering the fact that all of the three main approaches outlined here appear to be committed to the idea that there can be collective forms of intentionality, it is just that the significance and basis of this claim varies between them.

having a mind. In other words, the realm of the *mental* is held to exist entirely within the confines of the individual's head. However, it argues that we can correctly ascribe thoughts and/or actions to the group as a whole in virtue of some of the peculiarities of this individualistic mental life. It is in this sense that what this approach aims to provide is, as Schmid refers to it (2009, 33), a sort of *domestication* of collectivity. This is to say, that the aim is to provide an account of thinking and acting together that enables us to correctly speak of a group thinking or acting, while remaining compatible with the everyday intuitions that underpin individualism. The pertinent thought is that this might be achieved by admitting collectivity, in one way or another, into the *mode* or *content* of the intentional states that are had by individuals when acting as part of a group. To put it simply, and broadly, the argument is that we have individual subjects with individual minds, thinking in distinctively collective sorts of ways.

The second approach, most commonly referred to as *distributed cognition*, developed out of the cognitive sciences and directly challenges that which is accepted by both individualists and supporters of the intentionalist approach alike: that only individuals have minds. In general, this approach argues that the traditional conception of cognition as wholly contained within the brain, and associated with the phenomenally conscious individual, is mistaken. Instead, by reconsidering the category of cognition we can understand it as naturally incorporating elements not just outside of the skin and bones of the individual, such as artefacts, but also, at least potentially, as extending to groups of individuals. More specifically, however, this approach has proceeded along two related, but quite distinct, lines of thought. The first adopts a functionalist conception of the mind and cognition, arguing that it is on this functional basis that we can conceive of elements outside of the skin and bones of the individual as forming part of the individual's cognitive system; a system that includes the brain but which should be understood as extending out from this cerebral starting point. Further, it is also argued that it is on this basis that we can think of groups of individuals as constituting, functionally speaking, a single cognitive system. This is a position that has developed out of the *extended mind hypothesis* fostered by Clark and Chalmers (1998), as well as seminal work by Hutchins (1996) and, more recently, Tollefsen (2006). The second takes an even more radical and *enactivist* approach, arguing that cognition must be understood not as taking place in the head in the first place, but as involving a dynamic relation between mind, body, and world. In other words, it blurs the lines that have traditionally been held to separate both the cognitive individual from the world in which they cognise and distinct individuals from one another's cognitive lives. This position has been heavily influenced by work from Noë (2006), Thompson (2010), and Hutto & Myin (2013). Regardless of the differences between these two strands, what they share in

common is the thought that by re-evaluating what is involved in thought and action, we discover that many of the objections to their being ascribed to groups disappear. In other words, individualism has just got it wrong about what is involved in thought and action in the first place.

The third approach, which will be referred to as *plural subject theory*, is rarely discussed in its own right, with the various positions that fall under it often being associated with one of the other two categories. However, these positions are united, and distinguished from the other two broad approaches, in virtue of their commitment to the idea that groups can think and act on the basis of forming a single subject of that thought or action, that is, a *plural subject*. Included in this approach is the work of the likes of Margaret Gilbert, Carol Rovane, and Michael Bacharach¹⁰. Undoubtedly there is a great deal that separates the work of these three authors. Most significantly, while Rovane and Bacharach argue that subjecthood¹¹ can be ascribed to groups on the basis of a re-evaluation of what individual subjecthood consists in, Gilbert appears to retain an intuitive conception of individual subjects, while arguing that groups can count as subjects in a quite different sense. However, they are all united by the thought that the reason that we can make sense of ontologically collective forms of thought and action is because to be a subject is, at the very least, to be able to think and act, and groups can be subjects in their own right.

In considering some of these ideas, individual criticisms will be made of at least two of these approaches¹² on their own terms. However, viewed from a bird's eye perspective, the aim of this thesis is to undercut a more fundamental assumption that has, until now, structured the way in which these approaches have been framed. Specifically, all of the respective theorists¹³ involved assume, in one way or another, that if we want to ascribe thought or action to groups then we will have to re-think either thought, action, or the thinking and acting subject. Either our conception of thought and/or action at the individual level will have to change (in the case of theories of distributed cognition and some versions of plural subject theory) or thinking and acting together

¹⁰ Bacharach might strike some as a surprising addition to this list, considering the fact that his work is rarely, if ever, discussed in these terms. However, as will be shown in Chapter 4, there are significant reasons to take the book that he was working on at the time of his death, *Beyond Individual Choice* (2006), to contain a significant argument in favour of the possibility of plural subjects.

¹¹ Throughout this thesis the term *subjecthood* will be utilised to refer to the quality of being a subject. This is because in the context of many of the arguments considered below, the term *subjectivity* would be ambiguous between a form of self-awareness and the object of that awareness.

¹² While the intentionalist approach will be considered in detail in Chapter 2, and plural subject theory will be given a sustained analysis in Chapter 4, explicit reference to theories of distributed cognition will be largely absent from this thesis. This is partly because such a consideration will not prove necessary to developing the account offered below. But, more importantly, it is because these positions bear a complex relation to the work of Merleau-Ponty, whose own thought will prove central to reaching the desired conclusion. As such, an independent consideration of distributed cognition would unnecessarily complicate matters.

¹³ As will be seen, the work of Hans Bernhard Schmid (2014b) is a notable exception.

will have to be understood as differing in important respects from the individual case (in the case of the intentionalist approach and other versions of plural subject theory). Of course, to a certain extent the group must be understood as differing in various respects from the individual, if only for the fact that the former are composed of the latter. However, part of the motivation behind this assumption is a further assumption, held to be absolutely indubitable by all concerned, that the traditional view of the individual as a phenomenally self-aware¹⁴ subject or agent is a state of affairs that cannot possibly hold in the case of groups. In other words, if thought and action can only be attributed to a phenomenally self-aware subject, then the idea of ontologically collective forms of thought and action cannot even get off the ground.

It is difficult to pinpoint the precise source of this, perfectly reasonable, view. It is perhaps founded on a belief that phenomenal self-awareness is intrinsically tied to a brain. Thus, while intentionality, cognition, and subjecthood might be features that we can conceptualise as not necessarily restricted to the confines of a single central nervous system, phenomenal self-awareness is not open to such re-conceptualisation. As such, if we want to ascribe properties like intentionality, cognition, and subjecthood to groups, we will have to construe these in such a way that they do not require phenomenal self-awareness. For, what a group certainly does not have is a single unified brain.

Nevertheless, the aim of this thesis is to convince the reader that precisely this seemingly absurd possibility is just that – a possibility; it is not only conceivable but perfectly reasonable that groups can obtain a form of phenomenal self-awareness, and it is precisely this *plural self-awareness* that constitutes them both as a subject and as a minded agent that can act in the world. To put it differently, it will be argued that it is possible for a group to think and act for almost the exact same reason that it is possible for an individual to think and act, that is, because it is a phenomenally self-aware subject. There is a sense, therefore, in which the aim is to defend what amounts to a particularly radical form of plural subject theory, albeit one which takes seriously that which is dismissed by current proponents, and to thereby convince the reader of something that has been discarded almost entirely from the contemporary attack on individualism, as an idea that belongs in the nineteenth century when it was last taken seriously. Most importantly, however, the aim is to do so in a way that is continuous with the reader's intuitions regarding the nature of our social experiences. What we will be left with is a conclusion that, while it seems radical at first, is based on entirely reasonable premises and which is able to make sense of our actual experiences of

¹⁴ While the following will speak for the most part in terms of awareness and self-awareness, it will be assumed that, in the present context, that these terms are synonymous with consciousness and self-consciousness, and will occasionally be used interchangeably.

thinking and acting as part of a group on an everyday basis. In other words, the aim is to render the seemingly radical and unintuitive, convincing and wholly intuitive.

IV: Phenomenology and the Philosophy of Mind

Convincing the reader of this possibility will clearly require that we examine in some detail exactly what is involved in phenomenal self-awareness at even the most basic and fundamental level. In turn this will require that we come to grips with that area of philosophy that deals with the nature of self-awareness at this level in the greatest detail, namely, phenomenology.

In particular, it is possible to view the present thesis as structured around two key phenomenological concepts. On the one hand, an appeal will be made to an idea that is perhaps present in one form or another in the work of all the major phenomenologists, but has been articulated most clearly in recent years by Dan Zahavi. This is the idea that what it means to be a subject, at least in the most minimal or fundamental sense, is to be *pre-reflectively* self-aware. More specifically, the idea is that all experience must be understood as intentional not only in the familiar sense that it is presented as *of-something* but also in the sense that it is presented as *for-someone*. It is the latter, the subjective character of experience, that Zahavi argues is to be understood as constitutive of subjecthood. The self-awareness in question is to be understood as pre-reflective in the sense that it is neither thematic nor an awareness of self in terms of certain definite properties. For, such a concrete or reflective form of self-awareness can only be understood as arising on the basis of this prior, and more basic, form of self-awareness.

On the other hand, the argument will also lean heavily on the work of Maurice Merleau-Ponty in order to develop the idea that this most basic pre-reflective sense of self must be understood as an intrinsically *bodily* form of self-awareness. This is not an idea that is peculiar to Merleau-Ponty among phenomenologists, but it is he who radicalises it and demonstrates the precise extent to which human thought is enveloped in our bodily nature. For Merleau-Ponty, to experience is to be given to oneself as a bodily being, both in and towards the world that is experienced, in a way that breaks down many of the traditional boundaries thought to hold between the subject, the object of their experience, and the world in which this experience takes place. It is these particular insights regarding the bodily nature of the subject and experience that will prove crucial to an understanding of how pre-reflective forms of self-awareness might be open to pluralisation.

Both of these ideas, therefore, will be defended in some detail (in Chapters 3 and 5 respectively) and these defences will form the fundamental building blocks out of which an argument in favour of the possibility of phenomenally self-aware groups will be formed. However, what is particular

about the way in which these concepts will be defended is that these defences will progress via parallel and complementary ideas that are found in what is generally considered to be analytic philosophy of mind. This move, from ideas that have been mainly dealt with in analytic philosophy to a phenomenological, and therefore continental, understanding of them, may strike many as controversial. Until relatively recently it was almost universally accepted that while phenomenology and analytic philosophy of mind might be broadly interested in the same sorts of things, they were also quite clearly distinct enterprises, separated by differing core assumptions about how we ought to set about the task of trying to understand the mind and experience. While no-one was able to say precisely what this division consisted of, most had a reasonably good grasp of what fell into each category. In other words, whatever the divide consisted of, most agreed that there was one.

By contrast, in recent years this universal acceptance has started to falter and many have challenged the idea that there is any clear-cut or obvious distinction between analytic and continental pursuits. In fact, it has become almost cliché to declare that one does not understand or recognise a distinction between the two. However, in the case of this thesis, the cliché is apt. The claim that is made, by pursuing the cross-disciplinary argument presented in what follows, is not that there is no such distinction to be made, for it is often quite clear when one is doing continental or analytic philosophy. Rather, the claim is merely that with certain subject matter there is no particularly good reason not to cross the boundary in question and the subject matter of this thesis is precisely such an occasion.

The following is neither a piece of phenomenology nor analytic philosophy of mind, and the argument as a whole should not be viewed as aimed at a particular tradition. While the conclusions that will be reached are largely phenomenological in character, they represent an attempt to contribute to a debate that has arisen primarily within analytic philosophy over the last 20 to 30 years. As such, the arguments in their favour should be taken to be inadequate if they fail to convince, or at least appeal to the intuitions of, the average analytically minded philosopher. In other words, it ought to be convincing across the philosophical spectrum and if either the staunchly analytic or the strictly continental philosopher can find nothing in it that would persuade them of the possibility of plural self-awareness, then it ought to be judged as unsuccessful.

V: Overview

With this modest aim in mind, Chapter 2 will turn to the first of the three contemporary forms of anti-individualism outlined above, that is, the intentionalist approach. In many ways this might look

like an unnecessary step. After all, if the aim is to undercut the assumption on which the work of those involved in all three of these approaches has thus far been based, by arguing that groups can achieve a form of self-awareness, why bother with a detailed exposition of any of them? However, in the case of the intentionalist approach there are at least two significant reasons to subject it to an in-depth analysis. The first is that it has, perhaps more than either of the other two approaches, played a vital role in the resurgence of interest in forms of anti-individualism. There is a sense in which it epitomises the thrust of this contemporary approach as a whole: the attempt to start from broadly individualist assumptions and still be able to defend a broadly anti-individualist view of thoughts and actions, or, as it was put above, a *domestication* of collectivity. The second is that the problems that are inherent to this attempt to domesticate the collective will structure the explication of an alternative; any satisfying defence of the possibility of ontologically collective thoughts or actions will have to be able to overcome those problems that the intentionalist approach is unable to. Specifically, what will be found is that this attempt to start from broadly individualistic assumptions leaves the intentionalist approach unable to account for the *plural-yet-unified* character of our experiences of thinking and acting together. Any alternative, therefore, will need to be able to accommodate the fact that we experience thinking and acting together as occurring at a level between the individual and the *supra-individual*, that is, at the level at which there is simultaneously both unity and plurality.

Perhaps a third reason to consider the intentionalist approach is that it is suggestive of the idea that, if it is not possible to account for ontologically collective forms of thought and action in terms of individuals thinking in particularly collective ways, then it is the assumption that there is only individual subjecthood that ought to be challenged. Before turning to current approaches to plural subjecthood, however, the question of what subjecthood consists of must first be considered. It is pointless to attempt to provide a defence of plural subjects if we are not first clear about what makes an individual a subject. After all, this is a topic about which there is very little agreement in its own right. Partly in virtue of this lack of agreement, there will be no attempt to provide a full or complete account of subjecthood, but instead, in Chapter 3, an attempt will be made to present what amounts to a partial or minimal account. By pursuing a, slightly unusual, line of thought that will proceed through the work of Sydney Shoemaker and Immanuel Kant, before ending in the recent phenomenological ideas of Dan Zahavi, a position will be defended that argues that to be a subject, at the most basic or fundamental level, is to be given to oneself as a subject by experience,

that is, to be *subjectively self-aware*¹⁵. As Zahavi himself puts it, the idea is that subjecthood is constituted by the pre-reflective first-personal character of all experience; the *for-me-ness* of experience.

Thus, if to be a subject is to be subjectively self-aware, this would suggest that to be a plural subject would be to possess a form of *plural* subjective self-awareness. However, the problem that this presents for an account of plural subjecthood is that, as suggested above, it is precisely ideas of this sort that contemporary forms of anti-individualism were trying to avoid. With this in mind, in Chapter 4, four different attempts to defend the possibility of plural subjecthood will be considered. The first three, offered by Gilbert, Rovane, and Bacharach, all attempt to do so while avoiding the idea that groups can be phenomenally self-aware subjects in their own right. Nevertheless, it will be argued that all three fail directly in virtue of this fact and, as such, legitimise a serious consideration of the alternative. The possibility of this alternative, that there might be forms of plural subjective self-awareness, has been recently defended by Hans Bernhard Schmid. However, while Schmid's account offers notable improvements on the previous three accounts of plural subjecthood, it will be shown to give rise to a problem that does not just call into question the possibility of plural subjective self-awareness, but also exposes a fundamental flaw in the account of subjective self-awareness outlined in Chapter 3. Specifically, what Schmid's argument unwittingly brings to light is that this account has a problem of individuation. This is to say, that it is not clear how subjective self-awareness could constitute a particular individual subject, let alone how a plural form of subjective self-awareness could be constitutive of a particular plural subject.

Consequently, with this initial foray into the idea of plural subjecthood ending in failure, in Chapter 5 a retreat will be made to the territory of individual subjecthood in order to repair the account developed in Chapter 3, in light of the problems uncovered in Chapter 4. The suggestion will be that the necessary individuation of the subject can be accounted for in terms of the intrinsically bodily nature of subjective self-awareness. To put it differently, it will be argued that to be subjectively self-aware is to be given pre-reflectively to oneself in experience as the bodily subject of that experience, where such *bodiliness* must be understood as particularising and, as such, constitutive of an individual subject. Demonstrating the veracity of this idea, which amounts to a radical rejection of dualism, will involve grappling with concepts of embodiment present in both

¹⁵ This term is certainly not ideal, as 'subjective' always carries with it the unfortunate connotation of there being 'no fact of the matter'. However, considering the number of different ways this idea has been discussed under as many different terms, it is necessary to appeal to a generic term that allows us to avoid adopting the particular terminology of any one account in particular. For instance, Zahavi's notion of pre-reflective self-awareness is here taken as one particular version of the concept of subjective self-awareness.

contemporary philosophy of mind and phenomenology. However, ultimately it will be suggested that it is only the work of Merleau-Ponty, in virtue of the particularities of his phenomenology of perceptual experience, which can properly account for the bodily nature of the subject in a way that harmonises with Zahavi's conception of subjective self-awareness.

Finally, in Chapter 6, this Merleau-Pontian inspired account will be utilised and extended upon in order to revisit the idea of plural subjective self-awareness. Specifically, by incorporating some of Merleau-Ponty's thoughts on *habit* and *intersubjectivity* it will be found that, while Merleau-Ponty himself does not consider the idea of plural subjective self-awareness, he does nonetheless provide us with all of the theoretical tools necessary to construct just such a position ourselves. The suggestion being that collective forms of self-awareness can be made sense of in terms of groups of individuals partaking in processes of habitual familiarisation and thereby gradually building both a sense of shared experience and themselves as the single bodily subject of that experience. Moreover, it is these process of habitual familiarisation that can capture the plural-yet-unified character of our experiences of thinking and acting together that was so crucially lost by the intentionalist approach. Plural subjects, on this account, are to be understood as arising from the processes of development whereby a group, understood as a distributive set of individual subjects, gradually becomes supra-individual; a level of subjecthood that they can never actually achieve but towards which they are always moving.

With this overview in mind, it is worth noting in advance the double-edged nature of this thesis. While the explicit aim is to defend the possibility of ontologically collective forms of thought and action through an appeal to plural subjects constituted by a form of plural subjective self-awareness, and while the thesis is structured around this aim, there is a sense in which there is also an additional argument at work beneath the surface that relates to the nature of subjecthood in its most basic form. This additional argument begins in Chapter 3 with the defence of Zahavi's argument that minimal subjecthood is constituted by the pre-reflective first-personal presentation of experience. However, over the course of various points of critique and friction this account is supplemented, such that by the end what we will be presented with is an understanding of the minimal subject as temporally extended, inherently bodily, necessarily related to a more concrete or objective sense of self, and naturally extended beyond the confines of the skin and bones. What will be defended, in other words, is an account that understands even the most minimal subject as rich and complex in a variety of ways. This is worth bearing in mind simply because the success of this thesis as a whole rests on the success of this additional argument; it is part of the thought of this thesis that the

idea of plural subjective self-awareness only becomes a genuine possibility on the basis of this rich conception of the minimal subject.

Chapter 2

The Intentionalist Approach

Speaking in terms of an intentionalist *approach* is actually somewhat of a misnomer. Firstly, because it represents not one but two distinct, albeit related, positions, which will be referred to below as the *content* and *mode* accounts. Secondly, because there are few, if any, who have discussed these two positions in terms of constituting a single approach. Sometimes they are discussed in wholly distinct terms, but more often than not they are lumped together with what is described above as plural subject theory. Thus, for instance, in Schweikard and Schmid's *Stanford Encyclopaedia of Philosophy* entry on 'Collective Intentionality' (2013), they suggest that the field can be understood as being made up of three broad types of account, delineated by the element of intentional mental states that the relevant theorist holds to be open to collectivisation: the content, mode, or *subject*.

However, this is a misguided approach. For a start, and as was already suggested above, there is reason to believe that Collective Intentionality, as a field of research, has now become so broad that defining it in terms of just these three types of account is quite restrictive. Doing so closes it off from other avenues of contemporary anti-individualist thought that, in their own way, are perfectly happy to attribute intentional states to groups in one way or another. As such, there is a sense in which Schweikard and Schmid's categorisation is too restrictive.

Conversely, there is also a sense in which it is too inclusive, for it is surely odd to take the subject of an intentional state to be an *element* of that intentional state. If we take the statement of an intentional state such as 'You believe that it is raining outside' or 'I intend to go to the shops in the morning', it appears easy enough to analyse in terms of three constitutive parts. On the one hand we have the content of the state, or its intentional object, which is just whatever it happens to be about. In the case of the belief it is simply 'that it is raining outside' and with the intention it is 'that I go to the shops in the morning'¹⁶. Next we have the mode of the state, which refers to the kind of

¹⁶ This is actually a simplification, as it is a matter of some debate as to what the precise content of intentions must be in order for them to perform the role that is required of them (for instance see Searle 1983; Anscombe 1983; Thalberg 1984). In contrast to my belief that it is raining outside, which is satisfied just in case it is in fact raining outside, the satisfaction of an intention appears to require not just that the object of my intention comes about (for instance, that I do end up going to the shops in the morning), but that it comes about *intentionally*. In other words, it seems that it is not sufficient that the event which I intend to happen does in fact happen. Rather, it is also necessary that it comes about in the right way. And such issues are only a small part of a much larger debate regarding the relation between intention and intentional action. However, for the purposes of this chapter and the problems that it aims to raise, we can operate with a fairly naïve view of the content of intentions.

intentional state that it is, that is, a state of believing, a state of intending, a state of desiring etc. Lastly we have the subject, which is whoever the state belongs to or is attributed to; in the case of the intentional states mentioned above, ‘You’ and ‘I’ respectively. However, just because the linguistic statement of an intentional state submits to this kind of analysis, does not mean that it is accurate to take the subject of the state as a constitutive element of that state itself. Rather, it seems more accurate to view the subject of an intentional state as a matter of ownership or, as will be shown later, a question of the presentation of the intentional state as a whole (for-whom it exists). Moreover, to be a subject in the first place is to be something that is already capable of possessing intentional states. After all, how could one be a subject and not harbour any attitudes regarding the world? Accordingly, and as Schmid acknowledges elsewhere (2014a, 4–5), if one successfully defends the idea that there are non-individual subjects, then it would straightforwardly be the case that there would be collectively held intentional states (i.e. the intentional states of that non-individual subject). In this respect subject based accounts are doing something radically different to either content or mode based accounts, which appeal to features intrinsic to the intentional states held by individual subjects and the relations between them. As such, if versions of plural subject theory are open to criticism, they will be open to criticism of a very different sort (see Chapter 4).

By contrast, there is something that content and mode based accounts have in common that distinguishes them from every other form of contemporary anti-individualism, in that they both hold that the key to a full and convincing understanding of thinking and acting together lies in the capacity of individual subjects with individual minds to possess intentions¹⁷ that, in one way or another, make essential reference to the group in which they are participating. In other words, the intentionalist approach is distinguished by the fact that it holds this possibility to rest on each of the individual members of such a group having *individually-held-yet-collectivistic intentions*: intentions with either collective content or a collective mode.

As will be seen, the extent to which each of these positions can be understood as anti-individualistic, that is, the extent to which they defend the possibility of ontologically collective forms of thought or action, is both variable and open to debate. On one end of the spectrum, Bratman describes his own position as a form of ‘augmented individualism’ (2014, 11), and appears determined to stay as close as possible to the tenets of individualism. Conversely, it is readily apparent that Searle conceives of his position as a thoroughgoing anti-individualism, but, as will be suggested below, it

¹⁷ For the rest of this chapter the discussion will focus specifically on intentions to act, as opposed to intentional states more broadly. This is because while all versions of the intentionalist approach discuss intentions to act, they do not all accept that their analyses can be extended to other intentional states, such as belief, and the criticisms that will be made relate to the structure of the accounts on offer, rather than the specific intentional states to which they might apply.

is not clear that the radical internalism of his position allows space for such an understanding. As such, the only position that is undeniably in the anti-individualist camp is Tuomela's, which represents a very clear attempt to defend the possibility of genuinely ontologically collective forms of thought and action. However, despite these ambiguities, it is precisely the shared project of all three that will be brought under scrutiny. Regardless of the extent to which each of them can be characterised as anti-individualistic, the argument will demonstrate that they are all either entirely confined to an individualistic understanding of thinking and acting together, or they are inherently question begging.

Specifically, the challenge that this chapter will present is whether or not the admittance of collectivity into the individual's intending can tell us anything about the nature of the collectivity in question and, as will be seen, the answer is that it cannot. The idea that we can bootstrap collectivity off the ground, as it were, by placing that collectivity within the thoughts of individual group members is inherently flawed. The line of critique that will be pursued can be captured in terms of the following three claims: (1) both content and mode based accounts make references to 'we' a central feature of their analyses (sections I and II); (2) in the absence of an alternative, such references are ambiguous between a *distributive* and a *supra-individualistic* interpretation, and neither is sufficient to capture the *plural-yet-unified* character of our experiences of thinking and acting together (section III); (3) as such, the intentionalist approach both begs the very question to which it was supposed to provide an answer and, in doing so, exposes a serious challenge of its own to the success of any defence of ontologically collective phenomena (section IV).

I: Collective Intentional Content

With this in mind let us turn to the first of the two types of account that make up the intentionalist approach. The work of Michael Bratman is the clearest attempt to present such a position¹⁸. As suggested, like all versions of the intentionalist approach, he does not want to claim that there is some literal plural subject that can possess a form of intentional agency; only individuals are subjects and only individuals have minds of their own. Nor does he, in contrast to mode based accounts, want to argue that it is the type of intention possessed by individuals that is different when it comes to thinking and acting together. Rather, Bratman argues that what he calls a *shared* intention is

¹⁸ Many have also placed the work of Raimo Tuomela in this category, such as Searle (1990, 4040–406), Abraham Roth (Roth 2013), and (somewhat debatably) Gold & Sugden (2007b, 106–109). However, this seems to have been based on a slightly wrongheaded interpretation that emphasised aspects of Tuomela's analysis that are very similar to Bratman's (discussed in more detail below), while failing to recognise crucial differences. It is also an interpretation that Tuomela himself has frequently denied (for instance see Tuomela 1995, chap. 9; 2005).

something that is born out of the relations that can hold between a set of individually held intentions which each have the group as part of their content:

‘[S]hared intention, as I understand it, is not an attitude in any mind. It is not an attitude in the mind of some fused agent, for there is no such mind; and it is not an attitude in the mind or minds of either or both participants. Rather, it is a state of affairs that consists primarily in attitudes (none of which are themselves the shared intention) of the participants and interrelations between those attitudes.’ (Bratman 1999, 122)

It is in acting on the basis of such a shared intention that the behaviour of those party to the intention can be understood as a case of doing something together. In other words, and to borrow Bratman’s own example, what distinguishes a case in which two people are simply painting the same house alongside one another (i.e. with little or no interest in the activity of the other) and that in which they can be thought of as painting the house *together* in some sense, is that in the latter case they are party to, and thereby act on the basis of, a single shared intention to do so (Bratman 2014, 9–10).

To break it down to its very barest bones, this state of affairs roughly consists in something like the following: *x* intends *that we J*, *y* intends *that we J*, and there are suitable interrelations between *x* and *y*’s intentions *that we J*. Where the suitable interrelations include, for instance, that there is mutual knowledge between *x* and *y* regarding their intentions and that they formed their intentions at least partly in light of this knowledge. As Bratman succinctly puts it, ‘Our shared intention to paint together involves your intentions that we paint, my intention that we paint, and relevant further contents, interrelations, and contexts.’ (2009, 42)

As such, Bratman’s account involves intentions that violate the intuition that one can only intend actions that one believes to be within one’s own personal control or, as Bratman puts it, that ‘it is always true that the *subject* of an intention is the *intended agent* of the intended activity.’ (Bratman 2014, 13) In other words, just as it seems that *I* cannot intend that *you* paint the house (I can merely desire that you do so or believe that you will), it must also be the case that *I* cannot intend that *we* paint the house.

However, Bratman takes his account to be a natural extension of his planning theory of individual agency (2014, 4), and it is in light of this that such intentions that extend beyond the purview of personal control can be made sense of. On the one hand, he accepts something like a *Humean* or *belief-desire* theory of motivation, in which action (or at least the intention to act) requires a combination of a desire for a certain state of affairs to obtain and a belief that stipulates the best way in which to bring this state of affairs about. On the other hand, however, he suggests that there is an element that is missing in this view of practical rationality, in that we do not react purely to desires in light of our beliefs, but also in light of the partial plans that we build up over time. These

plans place extra demands for consistency on our processes of reasoning; when deciding upon what course of action to take I must not just take into account what I believe about the world, but also my existing plans (or you might say broader intentions) into which that action must fit (1999, 15–34). In other words, our decisions to act do not fit into a neat picture in which at every moment we are faced with a virtually infinite buffet of competing life plans. Rather, we make our choices in the context of an intricate web of interrelated intentions and plans that are built up over time, and while subject to re-evaluation, nevertheless place constraints upon our decision making.

In the context of this view of practical rationality Bratman fleshes his account out with the idea of *intending-that* (1999, 114–116). This is the sense in which I might intend ‘that my family stays healthy’; an action that I can perhaps aim to try my hardest to bring about, even if it cannot be said to be under my direct control. One might be inclined to argue that this notion of intention simply changes the subject at hand and that such *intentions-that* must be understood as a completely distinct form of attitude to ordinary personal intentions. However, such an inclination fails to recognise that intentions of this sort must play a key role if we accept Bratman’s idea that human beings are planning agents. After all, even one’s egocentric plans, such as my plan to write my thesis tomorrow, necessarily relate to an agent over which I do not have direct control, namely, myself tomorrow (a fact made acutely clear by how often such plans do not come to fruition). As such, just as my intention that my family stays healthy involves an outcome over which I do not have direct personal control, neither does my intention to write my thesis tomorrow; in both cases there is a sense in which the intended agent (‘my family’ and ‘myself tomorrow’, respectively) differs from the subject of the intention (‘myself now’, in both instances). For Bratman, the individually held intentions ‘that we J’ involved in a shared intention are of this variety. Such intentions will be a problem if one accepts, as many do, a causal account of intentions in which the defining property of an intention is its ability to directly affect action. But Bratman, while he accepts that causality must play some role in the carrying out of intentions, is inclined against what he deems to be a far too simplistic account in which all intentional action must be caused by a corresponding intention (1987, 112–126). As such, he has no problem with accommodating the idea that one can individually intend a collective action along these lines within his larger conception of agency.

As suggested above, by viewing his account as a natural extension of individual planning agency, part of Bratman’s thought is that we do not need to move very far beyond pure individualism in order to find the tools necessary to fully understand some of the peculiar features of social reality with which he is concerned. It is the very fact that part of what it is for individuals to intend and act is for them to build complex planning structures that inherently make reference to both themselves

at future times and others, which allows individuals to naturally take part in interrelated intentional structures that result in forms of agency larger than the individual.

The extent to which this account of thinking and acting together is anti-individualistic is an issue that will be returned to in due course. The important thought for now, however, is that, for Bratman, it is in virtue of the participant individuals in the shared intention making the group (the ‘we’ in ‘I intend that we J’) part of the content of their respective intentions that, when suitably linked together, those intentions can constitute a shared intention and make the resulting action (the group’s acting on their shared intention) a case of acting *together*. There is, therefore, a sort of reflexive co-constitution in play. The group does not have to literally or ontologically predate the intentions of the individuals in terms of being a subject in its own right. Rather, it is the individuals making reference to the group, almost as if there was in fact a pre-existing plural subject, which makes the shared intention and resulting action that of the group as a whole. As Bratman puts it:

‘[T]he *shared-ness* or *joint-ness* of shared intention consists of relevant contents of the plan states of each and relevant interconnections and interdependencies between the planning psychologies of each, all in relevant contexts.’ (2014, 12)

This is to say, that the participants intend to act as if they were a single plural subject and in doing so bring into being a form of shared intention that can only be correctly attributed to the group as a whole. In short, Bratman’s account relies on all of the participant members of the group making reference to the group activity as a whole as part of the content of their individually held intentions to act (2009, 42) such that, once connected appropriately and in the right conditions, these individual intentions with collective content constitute a single shared intention that cannot be ascribed to any of the individual participants, but only to the participants taken together.

II: Collective Intentional Mode

While mode based accounts share Bratman’s intuition that there are only individual subjects and individual minds, they seek to provide an account of thinking and acting together not in terms of the content of the relevant individuals’ intentions, but in terms of the form of those intentions. The general thought is that individuals are capable of thinking in fundamentally non-individualistic ways such that the resulting intentions, while held by individuals in the sense that they form part of their individual cognitive apparatus, are not like ordinary individual intentions. To put it differently, while there are only individual subjects and individual minds, those individuals can hold intentions that are of a distinctly *we*-form.

The work of Raimo Tuomela and John Searle are the clearest examples of mode based accounts. Nevertheless, it ought to be acknowledged that this is a slightly uneasy pairing. While there is a sense in which there is a general approach that the two share in common, it is also clear that there are fundamental differences in the kind of theory that they are advancing. For a start, only Tuomela is explicit in the assertion that he is putting forward a theory based on intentional modes, while Searle does not talk about modes at all. Moreover, Tuomela's work is particularly prone to alternative interpretations, drawing as it does on elements that can be found in Bratman and Margaret Gilbert's positions¹⁹. However, it is fair to say that they both argue that individual subjects have the capacity to form distinctly collective types of intention and make this ability the central tenet in their respective theories. It is in this light that it is useful to discuss their work under the same heading.

To begin with Searle's position, he presents what is by far the most straightforward account within the general intentionalist approach. What differentiates Searle's work most starkly from other theorists in the field, and accounts for the simplicity of his position, is that his is a non-relational account of thinking and acting together. This is to say, that for Searle there is a sense in which the substance of such thought and action all occurs within the mind of the individual; thinking and acting together is simply a matter of individuals having intentions in their heads that they attribute to themselves not in the first person singular but in the first person plural (Searle 1997, 449). In other words, where Bratman posits that the relevant intentions must be of the form 'I intend that we J', and that these intentions need to be suitably interrelated in order to establish a shared intention, Searle simply suggests that thinking and acting together is a matter of the individual members of a group all having intentions of the form 'We intend to J'. Where, for Searle, intentions of this sort, which he refers to as *collective intentions*²⁰, are just a fundamental part of our mental vocabulary in the same way that ordinary individual intentions are.

It is important to emphasise that this is not to say that the subject of such collective intentions is 'we', at least if the *subject* of an intention is understood as a type of ontological entity. For, Searle is quite clear throughout his work that, ontologically speaking, there are only individual subjects with individual minds. What differentiates a collective intention from an individual intention, therefore, is the fact that the former are had in a distinctly non-individual way. As such, there is a sense in which it would be more accurate to formulate the collective intentions in question as 'I we-

¹⁹ For similarities with Gilbert's position, particularly with regards to the role that commitment plays in collectively intending, see Tuomela 2000, 40–45; 2005, 331.

²⁰ It should be noted that this term is used very differently by other authors.

intend to J' (which would require that ordinary intentions be formulated as '*I I-intend to J*'). Indeed, it is quite plausible that it is Searle's failure to recognise that this would be a more accurate description of what he is getting at, and to properly differentiate between the mode and subject of an intention (where only the former can be properly thought of as an element of that intention itself) that explains his failure to acknowledge that his position is of the mode based variety.

That Searle proffers such an account is perhaps largely informed by his commitment to a particularly strong internalism, which excludes the possibility of there being an intentional state that is to any extent constituted by a relation to anything external to the individual's mind. For Searle, there has to be a sense in which a brain in a vat could be capable of everything necessary for thinking together with others, even if in actuality the lack of other people means that they are not doing so. Thus, he does not posit an intermediary group-level intentional state in the vein of Bratman's shared intention (or, as will be seen below, Tuomela's *joint intention*). Rather, acting together is simply the outcome of the relevant members of the group all harbouring intentions with the same collective mode and these collective intentions bringing about each of their participatory actions in the appropriate way.

This latter point is crucial, for there is still a sense in which these collective intentions need to operate through ordinary I-intentions (that is, intentions held by individuals in an individual mode) in order to bring about actual action, because each of the participants needs to be moved to perform their part in the group action. However, this does not mean that the former is reducible to the latter (Searle 2007, 50–55; 2010, 50–55). For Searle, no set of I-intentions, no matter how they are linked to one another, can add up to a collective intention, because the former cannot capture the element of *togetherness* that he believes is intrinsic to the latter (Searle 1995, 24–25). Rather, it is simply that part of what it is to possess a collective intention is to be committed to forming the relevant I-intention that corresponds to one's own part in the intended group action.

Most importantly for Searle, it seems that being able to form intentions with this kind of collective mode is a form of pre-linguistic *ability*, in much the same way that he thinks being able to I-intend is. Being able to think in terms of the group, rather than only from one's own individual perspective, is just a basic capacity that humans share in common. As such, he criticises other anti-individualistic accounts that rely on linguistic practices, forcefully arguing that language use, as a cooperative practice, already presupposes an ability to form such intentions. We would not be able to use language appropriately if we did not collectively intend to do so (2010, 49–50)²¹. Moreover, it is

²¹ This point, however, is a particularly controversial one, to which we will return several times below.

such a basic form of ability, he thinks, that not only can it be ascribed to human infants, but we must also assume that it is possessed by a range of animals as well (1995, 23)²².

The problem for Searle's account, however, is that as soon as it is recognised that he must be putting forward a mode based account of thinking and acting together, and it is realised that this account is motivated by his prior commitment to internalism, it becomes inconceivable that it can make sense of the possibility of ontologically collective forms of thought or action, even though this is quite clearly what Searle is trying to achieve. To see why, consider that an obvious critique of Searle's position might be to argue that it hard to see how any genuinely ontologically collective thought or action is going to arise on the basis of individuals intending in particularly collective sorts of ways, if thinking and acting together is not dependent on any relation between the intentions of those taking part. In other words, it seems like internalism and ontological collectivity are mutually exclusive concepts. Now, if one fails to recognise the crucial differences between the subject and the mode of an intention, then it looks as if Searle has a potential response to such objections. This response is to point out that *the one who acts* is a matter of to whom the intentions that resulted in the action were attributed, and in the case of Searle's collective intentions they are attributed to the group not the individual. As such, all that is required for an ontologically collective action to occur is that each of the relevant individuals holds, and acts on the basis of, an intention with the same collective character of being attributed to the group as a whole. However, once one recognises the differences between the subject and the mode of an intention, it should be realised that this idea of *attribution* is importantly ambiguous and it is this ambiguity that masks a subtle slide in Searle's position. If the idea is that the intentions are attributed to the group in the sense that the group forms the subject of the intention, then this appears to be precisely what has already been denied by Searle: there are no such non-individual subjects. If the idea is that the intentions are held by

²² Unfortunately, when it comes to explaining what this ability consists in or what it entails, Searle famously says very little. For the most part he seems merely content to show that his idea of collectively intending does not conflict with the account of ordinary intentionality that he had established in his earlier work or the idea that all intentionality must exist inside someone's head (1990; 2010, 47). It is this lack of a concrete explanation of how such an ability is possible that has led many to the same critique of Searle's account: that it is, in some sense, empty and lacking an actual explanation of what collectively intending consists in (Fitzpatrick 2003; Meijers 2003; Zaibert 2003). For Searle, we just happen to have a lot of intentionality in our own individual minds that is in the mode of the first person plural rather than singular (Searle 2007, 10). One of the few things that he does say on the subject, and which he affirms repeatedly throughout his work, is that collectively intending, and the resulting behaviour, are *biologically primitive phenomena* (e.g. 1995, 24), and that they are grounded in a 'preintentional sense of "the other"' as a possible cooperative agent and a higher order 'sense of community' (1990, 413). But he says little else to flesh out these rather vague comments on the subject. It is also worth pointing out that Michael Tomasello, who appears to be equally committed to the idea that forming individually held intentions with a collective character is a biologically primitive phenomenon for humans, denies that it could be possessed by animals as Searle assumes. In fact, Tomasello is largely concerned to show how the appearance of the capacity for such intentions at a certain stage in infant development is what distinguishes humans from apes who, he argues, do not develop the same capacity (Michael Tomasello 2009; Tomasello et al. 2005; Tomasello and Rakoczy 2003).

individuals but attributed to the group in the sense that they are intentions with a collective mode, as was suggested above, then it is not clear how this is supposed to avoid the objection just discussed. This is to say, that if there are no external relations between such individually held intentions with a collective mode, then it appears that the resulting action will merely be the action of a series of individuals all behaving in a pro-group sort of way. Searle might be correct that thinking and acting together is a matter of individuals having intentions of this sort, but this is not a view of things that could possibly progress beyond individualism, as there are neither ontologically collective thoughts nor actions in such a case.

By contrast, while Tuomela presents a position that is still focused on the mode of the individual intentions in question, his rejection of internalism appears to put him in a better position to defend the possibility of ontologically collective forms of thought and action. In fact, there is a case to be made for thinking of Tuomela's position as an amalgamation of elements found in the work of Bratman and Searle (as well as Margaret Gilbert's work), which ought to be taken to be a theoretical strength rather than a weakness. With regard to Tuomela's view of the relevant intentions in question (which he calls *we-intentions*), the only difference between his and Bratman's position is that where Bratman argues that the intentions had by the participant individuals need to be of the form 'I intend that we J', Tuomela argues instead that they need to be of the form 'I intend to do my part in X', where 'X' is understood to be the goal of the group not only in the sense that it cannot be achieved alone or unilaterally, but in the sense in which it is fundamentally seen as *ours*. Just like Bratman, Tuomela describes these individually held intentions as ones with 'a "holistic" content involving joint action' (Tuomela 2000a, 41) and argues that they must be linked together with various conditions and contexts such as mutual belief (Tuomela and Miller 1988, 375). Moreover, Tuomela describes these *we-intentions* as the individual's 'slice' of the *joint intention*, where the latter roughly equates with Bratman's notion of shared intention (i.e. the overarching intention that can be attributed to the group as a whole) (2005, 330).

Beyond this basic structural similarity, however, it is clear that Tuomela's account differs in far more important respects from Bratman's. Namely, for Tuomela, if such a process is able to operate it is only because certain other features are already at work behind the scenes, as it were. Most importantly, Tuomela introduces the distinction between thinking in the *we-mode* and thinking in the *I-mode*, with the former holding the key to understanding how the relevant *we-intentions* held by the individual group members are already a distinct form of intending.

This idea of the *we-mode* places a large emphasis on the participant individuals' prior, and seemingly innate, ability to think and reason from the perspective of the group (and in this respect

bears a certain similarity with Searle's position). In other words, it is this prior ability that establishes we-mode intentionality and allows the participants to think in terms of the intentions of the group as a whole. Thus, when Tuomela argues that the we-intention is the individual's *slice* of the group-level joint intention, this is not because the joint intention is constructed out of such slices (in a similar vein to Bratman), but because said individual has already been able to reason and think from the perspective of the group, ascertain what the intention of the group is or should be, and then derive their individual part in said intention:

‘The formation of a joint intention (and hence we-intention, a personal “slice” of the joint intention) requires that the participants jointly make up their minds to [jointly see to] something’ (Tuomela 2005, 330)

In other words, being able to make up our minds as a group, in virtue of each of us being able to reason from the perspective of the whole, is prior to any intention formation. Each participant is able to make reference to the collective as part of the content of their individually held we-intentions because said intention is already itself derived from the joint intention; it is already suffused with the collective perspective²³ such that, as he puts it, ‘we-intentions reflect the idea of a group at the level of the individual’ (Tuomela and Miller 1988, 371)²⁴. He also repeatedly emphasises the idea that these processes presuppose a pre-analytic notion of joint intention (Tuomela 2005, 342–343), as if being able to think in collective terms is just a fundamental building block of the human mind.

The overall picture, therefore, is of individuals partaking in thought and reasoning from the perspective of the group, deriving from the outcome of these thought processes their individual parts in the prospective ontologically collective action, and subsequently acting on these individually held we-intentions, which are suitably related such that they actually bring into being the ontologically collective action in question. It is in this sense that Tuomela highlights the fact that the theory is conceptually top-down and ontologically bottom-up (2013, 342). For, while the truly ontologically collective action can only be said to arise out of the interrelations between individually held we-intentions, this is only the case in light of the fact that those individual intentions were themselves formed on the basis of a conceptually collective perspective. Where for Bratman a shared intention, and the action it results in, is a construction of individually held intentions with collective content, for Tuomela the relevant individually held intentions are already of a

²³ In fact, in his latest work, Tuomela even presents this idea as building on and extending Michael Bacharach's work in game theory on the theory of framing and the role that perspective taking plays in reasoning and coordination (2013, 179–213).

²⁴ At one point, Tuomela even suggests that we can understand we-intentions as having a plural intentional subject and an individual ontological subject (Tuomela 2005, 343).

fundamentally different sort, as they are formed on the basis of we-mode, as opposed to I-mode, thought.

As such, while Tuomela is committed to the idea that ontologically collective action does require specific relations to hold between the we-intentions of individuals, what links both Searle and Tuomela's positions is that the possibility of ontologically collective thoughts and actions is premised upon the ability of individuals to think, and therefore intend, in a collective mode. Where content based accounts are based on individuals thinking individualistically and simply making reference to the collective as part of the content of their intentions, mode based accounts are based on individuals already being able to think in collective terms before they ever come to the matter of forming intentions.

III: The Ambiguity of the First Person Plural

Even on the basis on these rather brief summaries of the various positions on offer, it should be quite clear that the intentionalist approach in general is heavily indebted to references to the first person plural. While they each do so for different reasons, both content and mode based accounts make such talk of 'we' a central part of the analysis. For Bratman, the fundamental building block of a group's shared intention is the participant individuals all holding intentions of the sort 'I intend that *we* J'. While for Tuomela, whose position it will be assumed represents the archetype of the mode based account²⁵, it is that all of the individuals intend in a special *we*-, as opposed to I-, mode. Regardless of the specific merits and weaknesses that one supposes of the frameworks offered by each account, it is undeniable that references to 'we' are common to all. In fact, it is fair to say that *we-ness* has become a central motif of anti-individualistic discussions of thinking and acting together in general.

It would, therefore, seem prudent to investigate the nature of such references and the theoretical work that they are supposed to be doing. However, if we take the statement of an action including the first person plural, such as 'we are lifting the piano', it should be readily apparent that this statement is ambiguous in precisely the same sense as 'Tim and Tom are lifting the piano' (see Chapter 1, section II). In other words it is not clear, when taken alone, whether uses of the first person plural are meant to be read as referring the to the group in question in either an ontologically distributive or a collective sense²⁶. As has been suggested, 'Tim and Tom are lifting the piano' can

²⁵ It will also be assumed that even if Searle is able to respond to the objection outlined above, that his position will fall prey to the same kinds of criticism as Tuomela's.

²⁶ Such a statement is also ambiguous in the linguistic sense as well, but, as noted above, it can be assumed that the present discussion only concerns linguistically collective thoughts and actions.

be taken to refer either to [Tim] and [Tom] as individual subjects performing their own individual parts in a single linguistically collective action or to [Tim and Tom] as the single fused subject of the action, and references to 'we' would appear to give rise to exactly the same kind of ambiguity. While neither the content nor mode theorists wants to argue that 'we' can form the subject of an intention or action in any ontological sense, the same ambiguity of reference is going to arise no matter what role such references play in an intentional state: when we use the first person plural, to what kind of group are we referring? It is often the case that this ambiguity goes unrecognised by the theorists considered above. As such, it will be necessary to perform a bit of theoretical reconstruction and consider how they might have intended these references to be interpreted in a way that will perform the theoretical work that is being asked of them.

To begin with the ontologically distributive usage of the first person plural, there are good reasons to believe that if either content or mode based accounts are appealing to uses of 'we' in this sense then, while they might offer perfectly coherent accounts of thinking and acting together, these accounts will not be anti-individualistic. This is to say, that if the general tactic is to appeal to individuals intending in collective ways by incorporating references to 'we' within the content or mode of those intentions, then the resulting action will only be ontologically collective in its own right if these references to 'we' are taken to refer to the group in an ontologically collective sense.

This is perhaps clearest with regards to mode based accounts, where the term 'we' seems to be being used to express just the thought that individuals are capable of thinking in ways that make use of ontologically collective concepts. The very point of Tuomela's account seems to be that thinking and acting together is a matter of the individual members of the group being able to adopt the perspective of the group as a whole and form intentions on that basis. Thus, what distinguishes a we-intention from an ordinary intention, for Tuomela, is precisely the fact that it is an intention that is suffused with the perspective of the group understood as a single unit and aims to bring about a goal that can only be understood as that of the group as a whole. Conversely, if the 'we' of Tuomela's we-intentions was intended in the ontologically distributive sense, then it would just appear to be another, albeit slightly strange, form of I-mode intention. It is quite clear that Tuomela views his account of thinking and acting together as explaining the possibility of ontologically collective forms of thought and action and, insofar as this is the case, it must equally be the case that to intend in the we-mode is to intend in a way that makes intrinsic use of ontologically collective concepts.

It is in this regard that interpreting Bratman's position is made more complicated because, in contrast to Tuomela, it is not particularly clear whether the 'we' that he places in the content of

the relevant intentions *that we J* is meant to be understood as the ontologically distributive or collective usage. The closest he comes to clarifying the matter is in the following passage:

‘Well, in basic cases this use of “we” will ... be distributed. But we can here also, without circularity, avail ourselves of a concept of a group. We can do this if that concept of a group does not itself bring with it the very idea of shared intentionality. I might intend, say, that those of us in this part of the park run toward the hot air balloon that has crashed. If this use of “we” (or “us”) does not bring with it the very idea of shared intentionality, there need be no circle.’ (2014, 41)

The suggestion appears to be, therefore, that it is sometimes used distributively and sometimes collectively, although it is not entirely clear what difference Bratman believes that this makes. Moreover, understanding the import of this suggestion is further complicated by the fact that Bratman fails to distinguish between the linguistic and the ontological levels of ambiguity, leaving the possibility open that the *concept of a group* that he refers to in the above passage is simply a linguistically collective, but ontologically distributive, concept.

Nevertheless, such exegetical questions need not be of great concern with regards to present purposes. While it is the case that either interpretation can find some support in what Bratman says, only one interpretation makes Bratman’s position anti-individualistic, namely, that in which the collective content of the relevant individual intentions is taken to refer to the group in an ontologically collective sense. If, on the other hand, when the relevant individuals refer to ‘we’ in the content of our respective intentions, they are taken to be using the term in the distributive sense, there wouldn’t appear to be a basis for anything like an ontologically collective action, for no such thing is being referred to. If ‘I intend that *we J*’ is just taken to mean that ‘I intend that *I J*, and that *you J* ... and so on’, then what we would be left with would appear to be a set of individuals with a slightly arrogant or dictatorial approach to the actions of one another, in that they each seem to believe that they can intend the action of the other. This is not to suggest that such an account is necessarily incoherent²⁷, but it is to suggest that it is not an anti-individualistic account of thinking and acting together and is, therefore, of no use when it comes to the matter of establishing the possibility of ontologically collective forms of thought or action.

As such, what would seem to be required by either a content or mode based account in order to establish a form of ontologically collective thought or action, is an ontologically collective use of

²⁷ Although one is inclined to believe that, in spelling out the underlying commitment in this way, such an understanding does seem to stretch Bratman’s notion of *intending-that* to breaking point. Indeed, if Bratman would be willing to accept the possibility of intending that someone else does something, it would seem that what we are actually talking about is an other-directed wish or desire masquerading as an intention. For a more thorough discussion of the limitations of this individualistic reading of Bratman’s account see Petersson (2007).

‘we’ that forms an essential part of the mode or content of the relevant individuals’ intentions. Nevertheless, once this possibility is considered in more detail, the problem that presents itself is that the idea of an ontologically collective sense of ‘we’ is itself significantly ambiguous, in that there is more than one way in which such a sense might be ontologically collective. One way, for example, in which such references to ‘we’ might be understood is as referring to a group that is *supra-individualistic* in nature. In which case ‘we’ could be taken to be operating in much the same way as ‘I’, in that it refers to a single, unified, and relatively self-enclosed entity. Where the ontologically distributive sense of ‘we’ understands it as in fact referring to a series of implicit individuals, the supra-individualistic sense does not construe the ‘we’ referred to as consisting in any further subsidiary parts. Just as when someone declares ‘I believe that x’ they are not required to give some further story as to whom ‘I’ refers, the supra-individualistic reading of ‘we’ takes it that there are cases in which it can be declared ‘we believe that x’, where no further story should or can be given as to whom ‘we’ refers.

The most immediate problem with this supra-individualistic sense is that such a use of ‘we’ seems to be an incredibly unlikely interpretation of the way in which we understand and refer to groups on an everyday basis. In fact, such a reference to ‘we’ does not seem to capture what we mean by ‘we’ at all, but is rather just a reference to one more, albeit spatially disparate, individual. A useful way to come to grips with this problem is to consider a rather unlikely, but frequently insightful, source, in the form of *Star Trek* and the aggressive race of alien(s) known as *The Borg*. In the *Star Trek* canon, The Borg operate from what is referred to as a *hive mind*: each of the separate organisms, or *drones*, does not have an independent cognitive system but are all controlled by the same central collective cognitive system. Curiously, The Borg usually introduce themselves, in a monotonous yet chilling voice, by saying ‘we are The Borg’. It is curious because it seems that it must be incorrect. Rather, the introduction ought to be ‘I am The Borg’, because while there is a plurality of physical bodies there only appears to be one mind and, therefore, only one subject. In other words, The Borg appear to exhibit such an extreme level of unity that they cease to be a plurality of anything and instead become a single individual in their own right (albeit an individual more spatially complex than the average biological organism). It seems, therefore, that when it comes to the matter of understanding how we use references to ‘we’, there is no ‘we’ without a plurality of distinct ‘I’s, and there is no such plurality of ‘I’s in The Borg.

Beyond the matter of whether or not we do use the term in this sense, perhaps the more pertinent problem with regards to present concerns is that, as humans, thinking and acting together is never anything like the experience of being a member of The Borg. There might well be cases of human

thought or action that drift closer to the Borg-like case. For instance, at least historically, crowd behaviour has often been associated with processes of de-individuation, whereby one is said to *lose oneself* in the totality. Although such a view of crowds has become increasingly unpopular and a better example might be Georg Simmel's rather quaint treatment of the lovers and the priest, where for the former 'there is only a single life that can be lived or viewed from two sides', and for the latter 'the clerical function entirely supersedes and absorbs his individual existence.' (Simmel 2011, 13) However, if these are genuine cases, and one suspects that they are still starkly different from what it might be like to be a member of The Borg, they must surely be the exception and not the rule.

There is a sense in which if the theories discussed above were relying on this supra-individualistic sense of 'we' (putting aside for the moment the question of whether or not this is an acceptable use of 'we' at all) then the task of developing an account of ontologically collective phenomena would be achieved. But it would be achieved at the expense of being able to provide an account that accords with our actual experiences of thinking and acting as part of a group. As was suggested with regards to the distributive use of 'we', the importance that both mode and content based accounts place on such references means that the kind of reference is going to largely determine the sense in which the resulting thought or action can be said to be ontologically collective. If the use of 'we' is ontologically distributive, then the resulting thought or action will be that of the group only in a distributive sense. As such, it seems that if the use of 'we' is supra-individualistic in nature then the thought or action will be, in a sense, too collective (too Borg-like). It will be non-distributively that of the group, but in exactly the same way that the thought and action of individuals is non-distributively their own, which seems far too strong to capture the ways in which we do actually think and act together.

A more convincing depiction of ontologically collective forms of thought and action would appear to have to cope with the *plural-yet-unified* character of our experiences of thinking and acting together. This is to say, that defending the possibility of forms of thinking or acting together that go beyond a mere summation of individual thoughts or actions requires, first and foremost, an explanation of how individuality might be maintained in spite of unity. To refer to our experiences of being a group member as plural-yet-unified is to recognise the fact that ontologically collective forms of thought or action, if possible, must be understood as somewhere on a spectrum, from the wholly individual at one end to the wholly supra-individual at the other. All forms of ontologically collective thought and action require that our individuality is taken up and transformed by its being a part of the collective—partly subsumed into the totality—but that individuality rarely, if ever,

disappears entirely. Explaining how thought and action can occur within this strange grey area between the individual and the supra-individual levels is the true challenge for any theory that asserts the existence of ontologically collective thought or actions. As such, if content or mode based accounts rely on uses of ‘we’ that refer to an ontologically collective group in a way that is consonant with our experiences, these must refer neither to a mere summation of individuals nor to a supra-individualistic collective entity, but to a *unified plurality of individuals*.

IV: The Problem with the Intentionalist Approach

Once this difficulty of interpreting references to ‘we’ is uncovered, and it is recognised that an account that hopes to defend the possibility of ontologically collective thought or action must come to terms with the plural-yet-unified character of our experiences of thinking and acting together, it is easy to spell out the problem that this presents for content and mode based versions of the intentionalist approach, at least in so far as they are committed to anti-individualism. For, if the idea is that ontologically collective forms of thought and action might arise on the basis of individual subjects making reference to ‘we’ as part of the content or mode of their respective intentions, then this would appear to take it for granted that such references can refer to an ontologically collective group. In a sense, what they would appear to assume is that the nature of ontologically collective thought and action is obvious; the task is simply to show how it might be possible and we can do so by placing ideas of the collective within the content or mode of individually held intentional states in the form of references to ‘we’. However, this assumption is false; all that is clear about the nature of ontologically collective thought and action is that such phenomena must be understood as at once plural and unified. Presenting a theory that can accommodate such a tension is the true challenge.

In the absence of a response to this problem the positions that make up the intentionalist approach are seemingly left, at least to a certain extent, theoretically empty. On the one hand, if the point of Bratman’s content based account is to defend the possibility of ontologically collective forms of thought and action (which is, admittedly, unclear), then it becomes vague, or at least undecided, what the nature of the *sharedness* of the shared intention that is brought into being by our individually held intentions ‘that we J’ is. To put it differently, what does it mean to intend to do something as a group in a way that must be understood neither as a summation of individual actions to a common end, nor as the action of a single supra-individualistic entity? On the other hand, with a mode based account like Tuomela’s we are told that individuals are capable of thinking in intrinsically collective ways, but seen in the current light this tells us nothing about what it means to do so; what does it mean to think in plural-yet-unified terms?

In order to explicate the heft of this criticism it is worth considering a potential line of defence. For, it might be suggested that the theories above have been taken to purport to show something that they do not, or need not, purport to show. Specifically, perhaps the *emptiness* of which they stand accused does not necessarily represent the weakness it has been taken to be. Rather, one way to think of content and mode based accounts is as simply offering frameworks for which a certain amount of *filling out* will be necessary. In this light, these theories do not so much represent descriptions of how ontologically collective forms of thought or action actually occur, as they do explanations of how such phenomena might at least be metaphysically possible. Taken as such, they do not need to do all of the theoretical work, as it were, and they do not need to capture the exact nature of our experiences of thinking and acting together. Instead, one way in which this might be captured is by additionally appealing to a grasp of the underlying cognitive mechanisms in play when taking part in such collective forms of thought and action. As such, all that is required is that such theories be supplemented and supported by empirical work being conducted in either psychology or the cognitive sciences²⁸.

If this is the line that is to be defended there are two immediate problems that would need to be overcome. For a start, it looks like an extremely generous reading of the theories in question. There is no suggestion from any of the authors considered that they are attempting to provide anything less than a full account of thinking and acting together, and none of them imply that what they are presenting is something like a framework to be completed by the work of others. Secondly, beyond the question of the intentions of the relevant authors, it seems that not all of the positions discussed can even conceivably be thought of in these terms. While many of the things that Searle claims, particularly the idea that being able to collectively intend is a *biologically primitive phenomenon* (e.g. 1995, 24), seem conducive to supplementation with a cognitive underpinning, Bratman's position seems to be fundamentally formulated at the level of representations. This is to say, that the very point of Bratman's approach is to explain how we might provide an account of thinking and acting together purely in terms of individuals making overt representational plans with others, which appears to explicitly deny the possibility that the important work might be going on below the representational surface, as it were. Pacherie and Dokic (2006) have even forcefully argued that this is precisely the reason to favour Bratman's account, as it allows us to account for 'thick' forms of joint action that they believe cannot be captured in non-representational terms.

²⁸ For a particularly thorough attempt to explain how a mode based account might enlighten interactive theories of cognition, see Gallotti and Frith (2013).

These two difficulties, nevertheless, are really an aside to a far more fundamental problem, in that such a line of defence misses the thrust of the criticism being made, which is that a full and convincing defence of the possibility of ontologically collective forms of thought and action will have to allow for the existence of such phenomena somewhere on a spectrum from the wholly distributive to the wholly supra-individual. The significance of this problem is that it is virtually identical to the difficulty that seemed to motivate the generation of the intentionalist approach, and perhaps contemporary forms of anti-individualism in general, in the first place. When these theories began to arise the task that befell them was almost always parsed in the following terms. On the one hand, the aim was to move beyond the strictly individualist doctrines that had dominated methodology in the social sciences for most of the twentieth century and had proven to be quite incapable of providing a useful basis for social enquiry. On the other hand, one did not want, in escaping this individualist extreme, to be forced to defend a position reminiscent of the strongly collectivist theories that had been popular in the late nineteenth century and which often strayed into suspiciously supra-individualistic territory. While these might not experience the same methodological difficulties as the individualist doctrine, the view was that they avoided doing so at the unacceptable expense of an ontologically sure footing.

So, in effect, the very project of the intentionalist approach can be defined as the attempt to find a middle-ground between individualism and particularly extreme forms of anti-individualism that seem to be committed to the existence of supra-individualistic agents. As Schmid eloquently summarises, the task is that of ‘reconciling the *unity of the action* on the one hand with the *plurality of agents* on the other.’ (Schmid 2009, 23) However, as we have now shown, precisely the same task befalls attempts to interpret references to ‘we’ in neither distributive nor supra-individualistic terms. In other words, if the aim of the intentionalist approach is to explain how there can be ontologically collective thoughts or actions in a sense that captures their plural-yet-unified character, then making references to ‘we’ a central feature of such an explanation is inherently problematic, considering the fact that the first person plural gives rise to exactly the same question, that is, how such references can be understood as designating something that is held to be at once plural and unified.

To put it differently, for the intentionalist approach to present a successful form of anti-individualism, it would have to provide a theory that can capture the plural-yet-unified character of our experiences of thinking and acting together. However, if either content or mode based accounts rely on references to ‘we’ doing some of the essential work, as the above analyses suggest, then it should become clear that all they have managed to do is bypass the important question. For,

references to ‘we’ present exactly the same kind of dilemma, namely, of how to understand them as capturing the plural-yet-unified character of that which is referred to. Thus, the intentionalist approach attempts to capture the desired middle-ground between individualism and extreme forms of anti-individualism by appealing to the character of individually held intentions and positing individual subjects with individual minds that are, in one sense or another, thinking in we-centric terms and thereby forming intentions that have a collective character. But these we-centric terms are precisely the black box that is left unopened and unanalysed; the idea might be that ‘we’ can think or act because the individual members of ‘we’ can think in terms of ‘we’, but what it means to think in terms of ‘we’ is exactly what is not explained. As such, the line of defence taken above, that content and mode based accounts offer frameworks rather than complete descriptions, is simply not available. The charge is not that these theories do not perform all of the theoretical work that is being asked of them, it is that they perform *none* of the theoretical work. The intentionalist approach, as it stands, gets us no closer to a solution to the problem that it itself set up.

This line of criticism shares important affinities with another, now reasonably familiar, attack on the intentionalist approach. This charge, made most starkly by Petersson (2007) and Schmid (2009, 22–25), is one of circularity. Namely, that with regards to content and mode based accounts, the analyses that they offer are meant to explain what thinking and acting together consist in, yet, in one way or another, the idea of togetherness already features as part of the analysis. In other words, collectivity is part of both the explanans and the explanandum. In many respects the argument that has been made above also asserts that content and mode based theories are indeed circular. Specifically, the claim would be that the nature of references to ‘we’ is precisely what the intentionalist approach is meant to shed light on and yet it is these very references that are left as an unanalysed component by all of the relevant theorists.

However, the challenge made above also goes beyond the charge of circularity in a significant way. For, usually it is the case with such a charge that a sufficient response consists in demonstrating that the circularity in question is not vicious. This appears to be exactly how, for instance, Tuomela has responded to the issue of circularity. As emphasised above, he claims that his position is conceptually top-down and ontologically bottom-up. This is to say, that while collectivity features in the explanation, in the form of references to ‘we’, as well as being that which is supposed to be explained, this is possible because the way in which the relevant group members utilise references to the group, understood as an ontologically collective entity, does not rely on there actually being any such literal entity; the individuals think in ontologically collective terms and thereby, in a sense, bring that ontological collectivity into being. Nevertheless, it should now be recognised that this

response is insufficient. The problem is not just that the account is circular, but that this circularity effectively masks that which is truly problematic about making sense of the idea of thoughts or actions being attributed to a group as a whole. Tuomela's assertion that the issue of circularity does not arise because the relevant individuals only have to think in ontologically collective terms, without there actually being such an ontologically collective subject or agent, is irrelevant once it is realised that the real problem is explaining what it even means to think in such terms. How can we understand such references as designating something that is at once plural and unified, and how can we defend the possibility of ontologically collective forms of thought or action while capturing the plural-yet-unified character of our experiences of thinking and acting together?

In Chapter 1 it was asserted that the task of defending the possibility of ontologically collective thoughts or actions was that of demonstrating how it is sometimes appropriate to view a plurality of individuals, party to a thought or action, as a single unit. This is to say, if 'Tim and Tom are lifting the piano' the question is whether there are cases when this should be interpreted as '[Tim and Tom] is lifting the piano', as opposed to '[Tim] and [Tom] are lifting the piano'. It must now be seen that this was itself a simplification, and that a full and convincing description of ontologically collective thought and action will require that we come to terms with the middle-ground between these two alternatives. Any explanatory theory will have to be able to accommodate the fact that our experiences of such phenomena range from the more or less distributive to those in which the individual could almost, if not wholly, be said to lose themselves to the collective.

V: Conclusions

Thus, the intentionalist approach leads us into something of a dead-end in terms of defending the possibility of ontologically collective forms of thought and action. Existing accounts all rely, through making ideas of 'we' central to their analyses, on some reference to the collective forming part of the content or mode of individually held intentional states. By doing so they are party to an assumption regarding the status of references to the first person plural, which on analysis turns out to be considerably problematic, in that it is not at all clear or obvious how to understand such references. Once we begin to analyse the first person plural in any detail we discover that it is suggestive of precisely the same middle-ground that the intentionalist approach set out to explicate. On the one hand we could understand references to 'we' in a distributive sense, but this would not allow either content or mode based accounts to move beyond individualism. Alternatively, we could understand them supra-individualistically, but this would not appear to capture the nature of our actual experiences of thinking and acting together (not to mention the fact that it would make

content and mode based accounts reminiscent of the extreme forms of anti-individualism that the intentionalist approach was meant to avoid).

Instead, in order to present a theory that both defends the possibility of ontologically collective thought or action and does justice to our experiences of thinking and acting together, such a position would need to be capable of explaining how ‘we’ can refer to something that is at once plural and unified. But this is exactly what remains to be explained by the intentionalist approach. It does not tell us how there can be thought or action that is simultaneously that of the individual members of the group as well as that of the group as a whole. It does not tell us how there can be thought or action that is neither ontologically distributive nor supra-individualistic.

The point of establishing this line of critique is not to simply discount one potential explanation of ontologically collective forms of thought and action on offer. Rather, the problems inherent in the intentionalist approach must be viewed as the main obstacle to any successful anti-individualist account. If what is sought is a defence of the possibility of ontologically collective forms of thought and action, it is the fact that our experiences of thinking and acting together are necessarily plural-yet-unified to which we must attend and which any convincing account must be capable of accommodating. Consequently, while this specific feature will not be explicitly considered again until Chapter 6, the ideas that will be developed in the intervening discussion must be thought of as largely guided by this most fundamental of problems.

It is partly in this vein that the following chapter will turn to the issue of subjecthood. Part of the problem with the intentionalist approach is that ‘we’ is a form of plural *self*-reference but it is taken for granted what such a self or subject would look like, were such a subject to exist. In other words, there is a general lack of recognition of the problems intrinsic to the very idea of a plural subject. Of course, content and mode based accounts do not want to posit that there is any such subject in a literal sense. However, what they do all rely upon is the thought that individual subjects are capable of forming ideas of, and making reference to, plural subjects while forming their own individual intentions, and what has now been shown is that even the idea of a such a subject appears to involve what looks like a paradox; the maintenance of a contradiction between plurality and unity. By exploring the issues and debates surrounding subjecthood simpliciter, it will not only be possible to come to terms with what it would mean for such subjecthood to be pluralised, but it will also be found that the intentionalist approach’s distrust of the idea that there might be plural subjects in a literal sense is unfounded.

Chapter 3

The Subject and Self-Awareness

In the previous chapter it was suggested that accounts that try to defend the possibility of ontologically collective forms of thought and action in terms of individually held intentions that make reference to the collective in either their mode or content are fatally flawed. In short they presuppose that which they are supposed to explain and are, as such, suggestive of the fact that it is the underlying view of subjecthood in play that needs to be dealt with. However, it goes almost without saying, that the question of what subjecthood consists of is one of the most keenly disputed areas of philosophy. In fact, there are few truly significant philosophers who have had absolutely nothing to say on the subject. As such, the purpose of this chapter is not to provide anything approaching a full response to this question. Nevertheless, the intention is to present at least a partial answer. This is to say, the aim is to elucidate what amounts to the most basic and fundamental criterion of subjecthood; a criterion that is both necessary for any form of subjecthood and sufficient for the most basic sense of subjecthood. While being a subject or self in the fullest or richest sense will amount to more than fulfilling this core requirement, there could be no subjects or selves in any sense if this criterion was not met first and, if it is met, then there are subjects in at least some sense.

The ideas that are central to the account in question have been discussed, in one form or another, in a surprisingly wide array of areas and, as will be seen, certainly find a voice in both analytic and continental traditions. This account, broadly understood, argues that to be a subject in this most minimal sense is to possess a certain form of self-awareness, namely, awareness of oneself as subject or, as it will be referred to here, *subjective self-awareness*. This might immediately appear to be a quite straightforwardly circular account: to be a subject is to be aware that one is a subject. Surely one would already have to be a subject for that subjecthood to be capable of being picked out in awareness? But the idea is that the sense of self-awareness being appealed to is more fundamental than the kind of awareness that we can have of ourselves as a determinate entity in the world or, as it will be referred to here, *objective self-awareness*. For, if the idea was to appeal to the latter, then the account would certainly appear to be circular, as it would place the possibility of being a certain kind of object on being aware that one is that kind of object. Rather, the point is that, in one way or another, one's minimal subjecthood is fundamentally tied to the fact that one is an experiential being in the first place. In other words, it is that one is a subject in virtue of the fact that one is given

to oneself as such in the very nature or character of one's intentional experiences of the world. It is, therefore, a position that ties the concept of subjecthood very closely to that of phenomenal self-awareness²⁹. Simply put, to be a subject is to be self-aware, and to be self-aware is to be a subject, at least minimally speaking. Again, while there may be more concrete or richer forms of selfhood than that which is given by this core subjective self-awareness, they must be understood as building on top of this more minimal and fundamental subjecthood.

In order to defend this idea, a line of thought will be followed that will begin with a consideration of Sydney Shoemaker's argument for a form of non-perceptual self-awareness (section I), will proceed through a discussion of how this idea is pre-figured by Kant's account of pure apperception (section II), before turning to a consideration of the more recent phenomenological work of Dan Zahavi and the idea of pre-reflective self-awareness (section III). This might appear to be a somewhat convoluted route to take, but the purpose of this particular exposition is to demonstrate that it is by taking the notion of pre-reflective experience seriously that only phenomenology can provide a firm basis for the ideas that are developed in the work of Shoemaker and Kant. In short, it will be argued that to be a subject is to be presented to oneself pre-reflectively as the subject of one's experience. Finally, it will be suggested (section IV) that not only is this pre-reflective self-presentation constitutive of subjecthood in general, but that it is also both necessary and sufficient for the unity of the subject over time. In other words, subjective self-awareness is the most basic condition of possibility for personal identity in general.

I: Shoemaker and Non-Perceptual Self-Awareness

The most useful place to begin is with a debate that has been ticking over in analytic philosophy for the best part of the last 50 or 60 years. This debate, most directly inaugurated by the work of Sydney Shoemaker, centres upon the peculiarities of self-reference and self-awareness. In the mid-sixties Shoemaker, drawing on a distinction that he attributed to Wittgenstein, highlighted two distinct uses of 'I' which in the *Blue Book* Wittgenstein referred to as the 'use as object' and the 'use as subject'. Most significantly, Shoemaker noted that the latter, or at least a statement that utilises such a self-reference, is 'immune to error through misidentification relative to the first-person pronoun', while the former is not (Shoemaker 1968, 556).

To illustrate this difference, consider the following statement that contains two distinct references to self:

²⁹ As above, 'awareness' will be understood as synonymous with consciousness in this context (see footnote 14).

‘I believe that I am the one in the mirror’

Specifically, we have an ‘I’ to which a certain belief is being ascribed and an ‘I’ that is seen in the mirror. Moreover, it seems that we have one ‘I’ that might be used erroneously due to a misidentification of the subject in question, while the other is immune to this sort of error. In the case of the ‘I’ that is seen in the mirror (‘...I am the one in the mirror’), it seems that the person in question could be wrongly identified. This is to say, I could be aware that I see the reflection of someone in the mirror, but wrongly take that person to be myself. For instance, perhaps I am mistaken to think it is a mirror at all, when in fact it is a pane of glass on the other side of which my doppelganger is standing and mimicking my movements. However, in the case of the ‘I’ to which a certain belief is attributed (‘I believe that...’) it seems that I could not possibly be mistaken in the same way. As Andrew Brook puts it, ‘in some situations we cannot become aware of a person by being aware of certain experiences, take that person to be oneself, and be wrong.’ (2001, 28) In the case of the attribution of the belief (that I see myself in the mirror), I cannot come to be aware via experience that someone has a belief, take that person to be myself, and be wrong. I cannot identify the subject of the belief as myself, for example, when, in fact, it is you that holds this belief³⁰.

The former way of referring to oneself was that which Wittgenstein was referring to as the *use as object*, while the latter relates to the *use as subject*, and they appear to denote two fundamentally different senses in which one can pick oneself out. There is of course an important question as to which specific kinds of self-attribution involve the use of ‘I’ as subject and, therefore, exhibit immunity to error through misidentification. While classic examples tend to centre on immediate perceptual experiences such as ‘I have a toothache’ or ‘I see a field of heather’, there has been a great deal of debate over more contentious examples such as ‘I am wiggling my toes’ or ‘I saw a field of heather’. But the bigger question, with regards to present purposes, is what the source of this distinction is and why any of our self-references exhibit this form of immunity. What, in other words, enables this peculiar use of ‘I’ such that it cannot fail to pick out the correct subject?

³⁰ This is not to say that I cannot be mistaken as to the content of the belief in question. This is not so obvious with the mirror example, but it is clearly the case for beliefs such as ‘I believe that murder is always wrong’ (I might think that I believe this when in fact, on further reflection, I realise that I believe that in certain cases murder is permissible). But the important point is that in such a case the error in question would not be due to misidentification. It is not that I thought that I believed something when in fact it was someone else who was doing the believing.

Shoemaker's own line of response to this question is to point out that it is, first and foremost, because the awareness of oneself instantiated in the use of 'I' as subject does not involve any form of identification:

'My use of the word 'I' as the subject of my statement is not due to my having identified as myself something of which I know, or believe, or wish to say, that the predicate of my statement applies to it.' (1968, 558)

In other words, the use of 'I' as subject cannot misidentify because it does not rest on any process of identification in the first place, unlike the use of 'I' as object with which such an identification clearly is taking place. When I say that it is me that I see in the mirror, this can only be in virtue of the fact that I have identified the reflected image as myself in terms of the particular set of physical properties with which I associate myself and in virtue of it conforming to my expectations regarding processes of reflection. If the reflected image was presented at an angle that did not match my expectations as to how it ought to appear in relation to where I am standing, or if the reflected image bore physical properties that differ sharply from how I take myself to look, then I would not identify the image as a reflection of myself. However, such a process of identification could not possibly be taking place when I attribute to myself the belief that it is me that I see in the mirror. It is not that I recognise the existence of *someone* possessing such a belief and then, in turn, identify the subject of that belief via a set of specific properties of the believer, physical or otherwise. Recognising myself in the mirror is a matter of knowing what I look like as well as having at least a rudimentary understanding of how mirrors work, but recognising that I am the subject of a certain belief would appear to have nothing to do with knowing how *I think like*, as it were.

Conversely, it can be said that these two distinct forms of self-reference relate to two different ways in which we can gain information about ourselves; two different *modes* of self-awareness. On the one hand, there is an objective mode of self-awareness, which relates to the awareness we have of ourselves as a particular entity in the world; an entity that is understood or grasped in terms of its properties and, therefore, open to identification (and misidentification) through those properties. On the other hand, there is also a subjective mode of self-awareness, which relates to the awareness of oneself we have that does not appear to go beyond the mere knowledge that one is a subject; an awareness that does not seem to require knowledge of any properties of that self in particular and, therefore, a form of self-awareness that enables a form of self-reference that does not rest on identification. Gareth Evans captures this latter point perfectly:

'What is it for Oedipus to realize that he is the slayer of Laius? One thing seems clear: it is not to realize that Φ is the slayer of Laius, for any descriptive concept Φ . It is not to realize that the son of Jocasta is the slayer of Laius, or that the man who answered the riddle of the

Sphinx is the slayer of Laius, because Oedipus might realize these things without realizing that he is the slayer of Laius (not knowing that, or having forgotten that, he is the son of Jocasta and or the man who answered the Sphinx's riddle); and he might realize that he is the slayer of Laius without realizing these things, for the same reason.' (Evans 2001, 96)

Shoemaker takes from this that we ought to give up on the idea that this subjective form of self-awareness can be captured in terms of a perceptual model. This is to say, that the awareness that one has of oneself as a subject is not a form of perception in which the self is taken as perceptual object. Whatever sense of self that one is aware of in subjective self-awareness, it is not a sense of self that can be known through its properties and hence it 'does not involve being presented to oneself as an object.' (Shoemaker 1968, 563)

However, Shoemaker also emphasises that we should not fall into the trap of thinking that, because the self of which we are aware in the subjective sense cannot be perceptually picked out, we have no awareness at all of the self in which the relevant properties are instantiated when we say things such as 'I have a toothache'. Shoemaker attributes this mistaken view to Hume when in *The Treatise* he famously claims that:

'For my part, when I enter most intimately into what I call *myself*, I always stumble on some particular perception or other, of heat or cold, light or shade, love or hatred, pain or pleasure. I never can catch *myself* at any time without a perception.' (1888, 252)

For Shoemaker, this epitomises the view that one can be aware of one's own perceptions but that this does not involve an awareness of the possessor of those perceptions, that is, an awareness of one's *self*; that, 'one is aware of the predicates of self-ascriptions, or aware of the instantiation of these predicates, without being aware of their subject, i.e., that in which they are instantiated.' (Shoemaker 1968, 563)³¹ Nor, Shoemaker argues, should we think that the 'I' in statements such as 'I have a toothache' is not doing any real work, but is rather a strange sort of placeholder. This is an idea that he attributes to Wittgenstein himself, who apparently was 'reported to have viewed with approval Lichtenberg's saying that instead of "I think" we ought to say "It thinks" (with 'it' used as it is in "It is snowing").' (1968, 555–556)³²

³¹ It is more than likely that this is an unfair representation of Hume's actual views regarding the self, which are frequently misrepresented. In fact, an argument could be made that his actual view is closer to that of the phenomenological account that will be developed in section III. This is an interpretation of Hume that is explicitly defended by Dorothée Legrand (2007a, 588). However, for present purposes it is unnecessary to deal with these complexities and Shoemaker's caricature of Hume's view is at least illustrative of the kind of approach that Shoemaker himself takes to be wrongheaded.

³² Interestingly, Gareth Evans (2001, 98) attributes an almost identical view to both Geach (1957) and Strawson (1964a).

Both Hume and Wittgenstein's mistaken views rest, for Shoemaker, on the same error. Both of them recognise the shortcomings of the perceptual model of self-awareness but fail to take this as a reason to give up on the perceptual model entirely. The result of which is incoherence:

'What makes the matter seem puzzling, I think, is that one starts off by trying to construe self-awareness on the model of the observational knowledge that a perceived thing has a certain sensory property, and fails to abandon this model, or to abandon it completely and consistently, when one becomes persuaded that self-awareness does not involve any sort of perception of one's self... One tries to construe one's knowledge of the instantiation of the attribute ascribed in a self-ascription on the model of a case in which one sees or otherwise observes the instantiation of a sensory attribute, like redness, while at the same time denying that one perceives that in which the attribute is instantiated. And this, of course, leads to incoherence. The way out of this incoherence is to abandon completely, not just in part, the perceptual model of self-knowledge.' (Shoemaker 1968, 564)

However, while Shoemaker might be clear that when it comes to subjective self-awareness that we ought to give up on the perceptual model, his characterisation fails to go beyond the negative. We have an awareness of oneself as subject that is non-perceptual and, therefore, leads to certain self-references being immune to error through misidentification, but how or why? Subjective self-awareness is not a perceptual form of awareness for Shoemaker, but what is the alternative?

II: Kant and Pure Apperception

One alternative is to think of subjective self-awareness in transcendental terms. Andrew Brook (2001) has advanced an interpretation of Kant that portrays him as not only aware of the distinction between subjective and objective forms of self-awareness, but also as having a positive account of subjective self-awareness that explains why it enables forms of self-reference that do not rest on identification³³.

Perhaps most significantly, for Kant, the two forms of self-awareness in question are explicitly distinguished in terms of the ways we have of gaining them. On the one hand objective self-awareness, or 'apperception in the empirical consciousness' (1999, A115), is an awareness of oneself that one might gain concerning the particular properties that can be ascribed to oneself; an awareness of oneself as a particular object in the world. It is, as such, an awareness of oneself that one gains through *intuition* of either *inner* or *outer sense*. The latter provides awareness of those properties that are, presumably, available for others to recognise, such as one's physical attributes. While the former provides awareness of those properties that are, for the most part at least, only

³³ The question of whether Brook's interpretation of Kant is fair will, for the most part, not be considered in any detail. Trying to simultaneously evaluate the argument that Brook perceives in Kant as well as whether or not this perception is accurate would be a huge undertaking. As such, other than at certain critical junctures where exegetical questions are unavoidable, the primary concern will be with the former.

available to the individual, that is, one's psychological states. In other words, the awareness of self in outer sense provides knowledge of those properties that one ascribes to oneself in virtue of them being the object of particular representations, while awareness of oneself in inner sense provides awareness of those particular representations themselves³⁴. But whether we are talking about psychological or physical attributes, the object of our representations or the very representations themselves, what is important is that the sense of self in question is the sense of oneself as the possessor of certain attributes and, therefore, as the particular entity that one is: identifiable, discoverable, locatable and, above all, individual.

On the other hand, subjective self-awareness is awareness of oneself gained through *pure* or *original apperception*. As Brook highlights, it is crucial to distinguish this sense of self-awareness from that gained through inner sense (2001, 17). Whatever pure apperception is an awareness of, it is not an awareness of the particular representations that we happen to have. Nor is it some peculiar extra representation of the self. Rather, and this is the crucial point as far as Brook's interpretation of Kant is concerned, we gain an awareness of ourselves through pure apperception by *doing* acts of representing:

‘How does awareness of our own acts of representing and of ourselves as their subject work? The act of representing makes us aware of three things. Consider the sentence:

1. I am looking at the words on the screen in front of me.

Kant's claim seems to be that the representation of the words on the screen is all the experience I need to be aware not just of the words and the screen but also of the act of seeing them and who is seeing them, namely, me. A single representation can do all three jobs. In Kant's words, the awareness of the latter two items is given “not indeed in, but with... intuitions”. (Brook 2001, 18–19)

In other words, there is the awareness of the objects represented (i.e. the words on the screen), but there is also the awareness both of the act of representing itself and, most importantly, myself as the subject of that act³⁵. Brook describes this as the *representational base* of self-awareness in the sense that it provides all of the representation necessary for all three forms of awareness.

³⁴ As Brook points out, Kant's account of inner sense is at best confused (2001, 17). As such, it is not clear whether or not Kant holds that in order to be aware of one's particular representations in inner sense, that these representations must themselves become the object of a further representation.

³⁵ This interpretation is certainly not uncontroversial. Most significantly, what Brook does not seem to even acknowledge are the frequent occasions when Kant describes the awareness of ourselves gained through pure apperception as a form of representation (for example at A401), albeit of a very specific kind that differs fundamentally from other forms of representation. Nevertheless, regardless of its exegetical accuracy, Brook's interpretation is interesting precisely because it is charitable. If Kant does hold that pure apperception is a form of representation, then it seems to add far more confusion to the position and makes it even more susceptible to the problems that are outlined below. In other words, there are issues with Kant's account even on Brook's interpretation, but they would only be emphasised on the basis of a more representation based interpretation.

The immediate upshot of this is that it renders Kant, quite rightly, a strong critic of *higher-order thought* theories, which argue that awareness of a representation (taken as the criterion for self-awareness) requires a further representation that makes the first its object (Brook 2001, 19). By contrast, for Kant, the awareness of self made possible through pure apperception is not the self as object of any representation. Rather, it must be distinct from the kind of awareness gained through intuition because that enabled empirical forms of self-awareness and, as Kant puts it, 'I cannot cognize as an object itself that which I must presuppose in order to cognize an object at all' (A402). Instead, the awareness of oneself that pure apperception provides is a very specific form of awareness. It is an awareness of oneself as a *transcendental subject*, that is, as the condition of possibility for thinking in general:

'Through this I, or He, or It (the thing), which thinks, nothing further is represented than a transcendental subject of thoughts = x, which is recognized only through the thoughts that are its predicates, and about which, in abstraction, we can never have even the least concept; because of which we therefore turn in a constant circle since we must always already avail ourselves of the representation of it at all times in order to judge anything about it; we cannot separate ourselves from this inconvenience, because the consciousness in itself is not even a representation distinguishing a particular object, but rather a form of representation in general, insofar as it is to be called a cognition; for of it alone can I say that through it I think anything.' (A346)

The transcendental subject, that of which we are aware through pure apperception, is a sense of subjecthood as merely the vehicle of representations; that 'form of representation in general' that renders representation possible.

Perhaps, however, what is most stark in the quote above is that this form of self-awareness is, for Kant, a form of self-awareness that does not involve the attribution of properties, just like that which Shoemaker believed made Wittgenstein's use of 'I' as subject possible. It is a 'wholly empty representation' of which we 'can never have even the least concept' when taken in abstraction from the particular thoughts that manifest it. As he crucially says elsewhere:

'[I]n attaching 'I' to our thoughts, we designate the subject... only transcendently, without noting in it any quality whatsoever – in fact, without knowing anything of it either directly or by inference.' (A355)

Brook therefore suggests that we describe this form of self-awareness as a *non-ascriptive* self-awareness; an awareness of self in which nothing is ascribed to the subject. The awareness one has of oneself as a transcendental subject, an awareness gained through performing acts of representing, is as a raw or abstracted subjecthood removed from the concrete condition of our empirical being. It is a subjecthood *as such* or, as Kant occasionally describes it (e.g. A355), an awareness only that we are a 'mere Something'. As such, where the empirical self, known through intuition, was

identifiable and individual, the self of pure apperception is blank, propertyless and, most importantly, universal. It is a sense of self 'which I can apply to every thinking subject' (A355). It is an abstract condition of the possibility of experience and, therefore, an awareness of nothing in particular other than that which is necessary to make sense of ourselves as a thinking being (and which I would need to posit in order to make sense of anything as a thinking being). So, to put the two into a hypothetical trans-historical conversation, it would seem that Kant would add to Shoemaker's assertion that subjective self-awareness is non-perceptual, that this is the case because it is transcendental; that it is a condition of the possibility of there being experience in the first place.

Moreover, and most significantly for our purposes, it should be noted at this point that, even though we have not discussed it in these terms so far, what Kant has presented us with (or at least what Brook interprets him as presenting) is the first formulation of the idea that subjective self-awareness (pure apperception) is constitutive of a minimal or basic subjecthood. The reason why is due to the nature of the subjecthood in question. The transcendental subject is not a subject in the sense of a concretely existing entity (which is to say it is not known through intuition). In fact, it is not a sense of subjecthood that goes, or could go, beyond the realm of thought. To think, on this account, is to be aware of oneself as a self, that is, as a transcendental subject as the possibility and ground of that thought. There is no transcendental subject, nor could there be, beyond the particular thoughts for which it acts as vehicle. To be a transcendental subject is only to think and to be given to oneself in thinking as the subject of that thought.

Regardless of whether or not Brook's interpretation of Kant is entirely accurate, it represents a view that is, in many ways, a very robust position and one which is quite close to the position that will be ultimately defended below. Nevertheless, there are significant, if quite subtle, problems with this account. They can be uncovered by asking a question of it that has been of great concern to many philosophers who have worked on the nature of and relation between mind, self, and consciousness: does all awareness entail a form of self-awareness? In other words, does Kant hold that intentional awareness (i.e. having a representation) is sufficient for, or necessarily involves, awareness of self?

On what has been said so far it might seem reasonable to assume that the answer to this question must be 'yes'. After all, according to Brook, Kant argues that a representation is all the representation necessary not just to be aware of the object of the representation and of the act of representing itself, but also of the transcendental subject of the representation. This would imply that thinking involves an awareness of self as the subject of thought and hence all transitive

awareness does entail a form of self-awareness. However, contrary to this intuition, for Brook the answer is a fairly emphatic ‘no’:

‘Is a global act of unifying representation also sufficient for being aware of oneself as subject? No; despite many claims to the contrary, notably by Strawson, Kant was clear that one could represent objects without being aware of oneself (A113; A117fn.; B132). If one can have representations of which one is not aware, as was suggested earlier, one could have global representations of which one was not aware – for example, if one’s attention was totally focused on the complex scene being represented. Even if each global representation is the full representational base of self-awareness, a direction of attention or some cognitive apparatus necessary for taking advantage of the available representational opportunities might be missing. (Something like this might explain why nonhuman animals are not aware of themselves as subjects.)’ (2001, 23)

To put it differently, while the representational base is necessary for awareness of oneself as the transcendental subject of one’s representation and all the *representation* that is required for awareness of self, it is still not sufficient for such awareness. For such awareness to manifest itself we must, in addition to having a representation, also attend to the act of representing.

Despite the fact that Brook seems to place very little weight on this feature of self-awareness (he does not mention it at all beyond the passage above), if this is a true depiction of Kant’s account then it gives rise to some serious concerns. In particular, on even a cursory reading it looks as if this account of self-awareness is going to be susceptible to a sceptical challenge that we have already considered. Namely, the one offered by Hume:

‘For my part, when I enter most intimately into what I call *myself*, I always stumble on some particular perception or other, of heat or cold, light or shade, love or hatred, pain or pleasure. I never can catch *myself* at any time without a perception.’ (1888, 252)

In other words, Hume’s concern is that when he performs some act of introspection, when he thinks about his thinking, all he can find are his particular thoughts. What he does not find is some extra thing in addition that he could call a self. But is that not exactly what Brook is suggesting Kant believes to happen when we attend to our own representations? That when I think about my thinking I become aware of myself as the transcendental subject of those thoughts; the ‘I think’ that is famously able to accompany all of my thoughts actually does; the subjectless representation ‘there is a tree over there’ becomes, upon attending to the act of representing, ‘I think there is a tree over there.’

Considering the fact that so much of what Kant had to say on many topics was a direct response to problems that Hume uncovered, it would be particularly odd to think that his account might be susceptible to such a basic concern of Hume’s. So let us consider a few ways in which the defender

of Kant's account might respond. Unfortunately, it is hard to see how any of them might be successful.

The first, and most obvious, response is simply to deny that Hume was right. In other words, it could be asserted that attending to one's representations does, in fact, yield awareness of self. Whether one finds this to be an adequate response is going to largely depend on one's intuitions, and there is very little that could be said here that would significantly alter these. For my own part I think that Hume's worry is a serious concern that must be dealt with by anyone wishing to develop a full and convincing account of subjecthood and, while this is speculative, I am compelled to think that Kant would also have thought so. Which is why I do not think that this is a serious line of defence for the Kantian to adopt.

A far more promising response would be to accept that there is some truth to Hume's concern but to suggest that his mistake was in searching for the wrong sense of self. In other words, the claim would be that the sense of self that Hume was looking for was the sense of self understood as an ontologically independent entity, but such a sense of self is not that which Kant associates with pure apperception. Rather, the self yielded through pure apperception is of a 'mere Something'; a sense of self only as the implicit subject of experience. As such, the Kantian can agree that attention to one's representations does not yield awareness of self understood as a concrete entity, but argue that it does yield awareness of self in a much thinner sense.

However, this only manages to sidestep the real issue. Regardless of the sense of self that Hume was looking for himself, the thrust of his concern should not be taken only to relate to such a particular sense of self and it is unlikely that it has been taken in this way by those who in more recent years have taken Hume's concern as a reason to be sceptical of the existence of a self³⁶. Rather, what Hume appears to be saying is that when he attends to his thoughts he does not find anything, that is, anything at all, in addition to those thoughts which might be associated with a subject or self. He finds no subject of his representations, only the representations themselves. Even if Kant is referring to a transcendental, as opposed to concrete, self, he is certainly committed to the idea that attending to one's representations yields fresh information; it yields an awareness of self as the transcendental and implicit subject of those representations. Moreover, upon attending to them, these raw representations are, necessarily, transformed in virtue of their being accompanied by an 'I think'. What Hume is surely denying is that any such transformation takes place.

³⁶ Again, it is important to note that it is also unlikely that Hume shared this scepticism (see footnote 31).

As such, one might be forced to turn to a third potential route out of this quandary, which is to argue that, at least on this point, Brook's interpretation of Kant is wrong. Or, perhaps better, that there is a more charitable interpretation of Kant that drops the necessity of this attention to one's representations for subjective self-awareness. Indeed, as Ewing points out, 'Kant wavers between maintaining that we are always conscious of [the transcendental unity of apperception] in some way and maintaining merely that we can become conscious of it at any time in connection with any representation.' (1996, 81) Indeed, despite Brook's lack of consideration of alternative interpretations, it appears to be one of those issues where Kant flagrantly contradicts himself.

The problem with this alternative interpretation is that it is left with a gap and without the necessary theoretical material with which to plug it. The problem is that there is a reason why Kant is forced to suggest that attention to one's representations is necessary to manifest awareness of oneself as transcendental subject. Because the transcendental subject is just that: transcendental. To say that something is transcendental is to say that it is a necessary condition for the possibility of something else. In this case the subject is the necessary condition for the possibility of experience. No subject means no experience. But to say that it is a necessary condition for experience is not to say that it is itself experienced.

As such, Kant has two options. The first is to argue that, while the subject and our awareness of it is transcendently presupposed by experience, it is itself not actually experienced in any sense; that an awareness of self as subject is never *manifested* in experience. In other words, he can deny that we ever have any experiential self-awareness (at least of ourselves as subject). However, this would involve a scepticism regarding the self that would strike most, including myself, as unpalatable. The second is to argue that such awareness of self is manifested by some extra faculty that allows us to obtain experiential self-awareness without having to resort to some further form of self-representation. In other words, this would appear to involve the ability to non-representationally attend to one's act of representing. Such an ability is clearly in need of detailed explication and it is not an explication that Kant himself, in harbouring a fundamentally representationalist view of the mind, is able to provide. However, considering the fact that the only alternatives are either scepticism regarding the self or some form of higher-order thought theory, it appears to be an option that ought to be taken seriously. As such, Kant's view of the self and experience presents us with a challenge: how can subjective self-awareness be given in experience without resorting to a representational account that looks like a form of higher-order thought?

III: Zahavi and Pre-Reflective Self-Awareness

In order to find an answer to this question it is necessary to turn to phenomenology. It is important not to understate the importance of this methodological switch, as it is precisely what is demanded by both Shoemaker's insights regarding subjective self-awareness and the inherent weaknesses in Kant's account. What the former clearly happened upon is the fact that subjective self-awareness, and the uses of 'I' that it enables, cannot be accounted for in terms of a perceptual understanding of awareness. As we saw, he argues that this model needed to be abandoned entirely. But while Kant offers us an account that goes much further in terms of trying to understand how subjective self-awareness is (non-perceptually) possible, it struggles to explain what it can mean for such awareness to manifest itself in experience. By taking the idea of *pre-reflective* experience seriously and by investigating the very nature of experience as it is presented to us, phenomenology offers the perfect solution to this problem.

Specifically, a view of self-awareness that has been articulated most clearly by Dan Zahavi in recent years, and which builds upon a line of reasoning that runs throughout phenomenological thought, is able to respond to many of the problems inherent in Kant's approach. Zahavi makes the same distinction between awareness of oneself as an object and awareness of oneself as a subject, but distinguishes between the two in terms of reflection. While objective self-awareness is associated with a reflective level of thought, subjective self-awareness is located at the pre-reflective level. Thus our objectifying sense of self, that is, our sense of self as a particular object in the world to which certain properties can be ascribed, is a result of taking a reflective stance on oneself; the result of holding some determinate feature of oneself as the object of an intentional state and thereby considering oneself in something like the same sense that others might consider one. By contrast, in Zahavi's view, subjective self-awareness is given in pre-reflective experience. It is a sense of self not as the object of any representation, nor is it a transcendental condition for the possibility of experience that is not necessarily itself experienced. Rather, it refers to the mode in which experiences pre-reflectively present themselves. In other words, the way in which they present themselves as *for-me*, not in the sense that they share the same referent ('me'), but in the sense that they share the same style of presentation, or qualitative *for-me-ness*³⁷:

‘The for-me-ness in question is not a quality like yellow, salty, or spongy. It doesn't refer to a specific content of experience, to a specific what, but to the unique mode of givenness or how of experience... It could consequently be claimed that anybody who denies the mineness

³⁷ From here on 'subjective self-awareness' is used exclusively to refer to Zahavi's notion of pre-reflective self-awareness.

or for-me-ness of experiences simply fails to recognize an essential constitutive aspect of experience.’ (Gallagher and Zahavi 2007, 50)

The heart of this idea is the simple observation that all experience is intrinsically, and necessarily, perspectival. There are no *free-floating* or *anonymous* experiences, but rather, ‘Every experience belongs to a subject, that is, either to me or somebody else; it cannot belong to nobody’ (Zahavi 2008, 46); they are all anchored to some perspective and angled toward a subject. For, it is not possible to make sense of the idea of an experience that does not have such a perspective. In many ways, therefore, the pre-reflective character of experience refers to just this perspectival structure or form, in which the subject is given as an almost implicit feature of experience.

Moreover, Zahavi is keen to emphasise that as a mode of presentation of experience, as opposed to an objectifying form of awareness, subjective self-awareness must be taken as ascribing nothing to the subject. The awareness of one’s subjecthood in question does not go beyond an awareness that one is a subject; that one’s experience is presented in a certain subjective mode. Indeed, he explicitly recognises the work of both Shoemaker and Brook, and agrees with them that subjective self-awareness must be understood as non-ascriptive (Zahavi 2007, 602; 2013, 332; Gallagher and Zahavi 2007, 55) and takes the idea of such self-awareness to be capable of enabling the peculiar forms of self-reference discussed by them.

Another way Zahavi’s particular take on subjective self-awareness can be understood is to consider that, in response to the question that was posed with respect to Kant, Zahavi would argue that all intentional awareness does necessarily involve a form of self-awareness. To be aware of something is to be pre-reflectively aware of oneself as the subject of that awareness, that is, as the perspective that makes that awareness a form of awareness. However, Zahavi is careful to distinguish this claim from other ways in which the claim is sometimes made. For instance, Harry Frankfurt affirms the claim that all consciousness entail a form of self-consciousness, but does so in the sense that he argues that consciousness is reflexively aware of itself (i.e. of that very consciousness), rather than aware of a self in the sense of the subject of that conscious state (Frankfurt 1988). Frankfurt denies that there is an awareness of oneself as subject in experience.

Zahavi does not entirely disagree with Frankfurt, in that he too thinks that there is a reflexive element to consciousness, but puts Frankfurt’s insistence that there is no awareness of oneself in consciousness down to his inability to distinguish between a more minimal and a more substantial concept of self. In other words, what Frankfurt is denying, like Hume, is that there is an awareness of self as a persisting and substantive entity, distinct from experience, in experience. Zahavi, by contrast, points to the fact that even though this is true there is still, and must be, a more minimal

awareness of oneself in experience; an awareness that does not go beyond the idea that experience is presented as from a certain perspective and as for-me:

‘The ego does not need to be conceived of as something standing apart from or above the experience, nor does one need to conceive the relation between ego and experience as an external relation of ownership. It is also possible to describe the first-personal givenness of an experience, that is, its very self-givenness or self-manifestation, as the most basic sense of self.’ (Zahavi 2008, 46–47)

It is in this respect that Zahavi’s account is able to respond to Hume’s scepticism regarding the self. It is not just that the awareness of self in question is minimal (after all, Kant also makes such an assertion), but that on this account subjective self-awareness is not something that we have to attend to for it to become part of our experience, and our attention would never be able to capture it fully anyway:

‘Although the anonymous functioning subjectivity cannot be thematized in its very anonymity, this does not prevent it from being given. Not only is it exactly characterized by its pre-reflective self-awareness, but we even encounter its elusiveness every time we try (and fail) to catch it in reflection, i.e., the reflection points toward that which both founds it and eludes it, and these features are not deficiencies to overcome, are not results that threaten the phenomenological enterprise, but are rather the defining traits of pre-reflective self-givenness.’ (Zahavi, 2002)

One could easily take this passage as a direct response to Hume. The idea being that Zahavi accepts Hume’s concern in the sense that reflection will never be able to capture the self as it appears in subjective self-awareness. Reflection will always, necessarily, miss its target because in reflecting it automatically changes the thing reflected upon. However, while Hume was right to suggest that on attending to his perceptions that he found nothing in addition to them that he might identify as himself, his mistake was in looking for something *additional* (a mistake that can also be attributed to Kant). For Zahavi, it is those very perceptions that are already intrinsically perspectival and subjectively given. One need look no further than these perceptions because they already are an acknowledgement of self in its engagement with the world or, to put it far more eloquently:

‘If we want to study the self, we should not look inside consciousness in the hope of finding some elusive I, rather we should look at our intentional experiences. Just as the self is what it is in its worldly relations, self-acquaintance is not something that takes place or occurs in separation from our living in a world. To put it differently, *our experiential life is world-related, and there is a presence of self when we are engaged with the world.*’ (Zahavi 2009, 563)

We do not discover the self in this abstract or minimal sense by, as so many philosophers have assumed to be the case, performing some special act of introspection. We discover the self, or better we discover our own subjecthood, in the way that our experience in interacting with the world is suffused by that subjecthood.

At the start of this chapter it was suggested that the aim was to elucidate the most basic or fundamental criterion of subjecthood, and there is good reason to believe that Zahavi's account of subjective self-awareness, understood in terms of the pre-reflective self-presentation of experience, satisfies this aim. To begin with, it appears that such self-awareness is going to form a necessary component of any form of subjecthood on this account. The relevant question, in this regard, is whether we can imagine there being such subjecthood, or indeed any sense of self, in the absence of there being first-personally presented experiences? And the answer must surely be 'no'. If experience lacked this character, if experiential life was, as Zahavi puts it, completely anonymous, then the very question of 'self' would not arise. Moreover, for Zahavi, the very fact that we are able to take ourselves as object demonstrates that we must already be given to ourselves as subject:

'It is this pre-reflective sense of self which provides the experiential grounding for any subsequent self-ascription, reflective appropriation, and thematic self-identification. Had our experiences been completely anonymous when originally lived through, any such subsequent appropriation would become inexplicable.' (Zahavi 2013, 334)

'To put it differently, a minimal or thin form of self-experience is a condition of possibility for the more articulated forms of conceptual self-consciousness that we incontestably enjoy from time to time.' (Zahavi 2009, 562)

With regards to this latter point the suggestion is not that subjective self-awareness is a 'condition of possibility' in the transcendental sense, but rather in the sense of being the raw material out of which more concrete forms of selfhood and objective self-awareness can be constructed. To know *oneself* as an object already presupposes that one knows *oneself* as subject, for otherwise the knowledge of *oneself* as object would not be of *oneself*, but knowledge of an object just like any other object. Thus, for ideas of subjecthood to even get off the ground, that is, for there to be subjects of any kind, it must be the case that there is experience that presents itself as for such a subject.

It might be contested, of course, that while such subjective self-awareness is a necessary condition of subjecthood in general, that it might not be sufficient. For instance, it might be suggested that some form of objective self-awareness is also a necessary component of subjecthood, properly speaking. On this view, the sufficient condition of subjecthood would be being both subjectively and objectively self-aware.

However, while it might be a largely empirical question whether or not there are in fact any subjects that do not possess a form of objective self-awareness, that subjective self-awareness is sufficient for a minimal subjecthood is shown by the fact that it is at least conceivable that there are. For instance, it is clear that non-human animals have experiences of the world and, on Zahavi's view, this entails that they will be pre-reflectively self-aware in the sense that their experience will have a first-

personal or perspectival presentation (for, they would be incapable of experience if this was not the case). But it is also at least plausible that they do not have any richer sense of self than this. This is to say, that many animals are perhaps unable to take themselves as the object of their awareness and thereby adopt a reflective stance towards themselves (in fact, it is plausible that this is what separates animals which are able to pass mirror recognition tasks from those which are not). Moreover, it is also plausible that Zahavi's position presents an intuitive way to understand the development of human infants. This is to say, that we might understand new-born babies as only being subjects in the most abstract sense, in that they are experiential beings, and as developing a sense of themselves as concrete entities through the process of maturation.

Thus, it is reasonable to think that subjective self-awareness not only presents a necessary condition for being a subject, but is also sufficient for a minimal sense of subjecthood. There is a sense, therefore, in which to be a subject, at least minimally speaking, is synonymous with being an experiencer; as long as there is experience there is the subject who is given as part of its mode of presentation.

IV: The Unity of the Self

So, on the view that has been derived from Zahavi's work, to be a subject, at least minimally speaking, is to be subjectively self-aware in a sense that is given pre-reflectively in the way in which experience is presented and which involves the ascription of no properties to that subject (precisely because such an ascription would involve the adoption of a reflective stance). Nevertheless, it might be objected at this point that this cannot be the whole story. While subjective self-awareness might ground the possibility of subjecthood, it appears to tell us nothing about the unity of that subject over time. It might be what makes me the subject of a particular experience, but is it what makes me the same subject over a series of experiences?³⁸ After all, the primary philosophical interest in the concept of subjecthood usually has to do with that which grounds the persistence of the subject. If an appeal to the idea of subjective self-awareness is unable to provide such a ground, it would appear to offer an account of subjecthood only in a very thin, and almost redundant, sense.

It is tempting to assume that the way in which to provide this ground is to move beyond subjective self-awareness and instead appeal to the relation between objective and subjective forms of self-awareness. For, awareness of oneself in the objective sense would seem to presuppose, at the very

³⁸ It should be acknowledged that the very idea of a temporally unified self has proven to be more controversial than the idea that there is some minimal awareness of oneself as subject contained in experience (for instance see Metzinger (2013)). However, this should be taken to be a sceptical position of last resort. For present purposes it will be assumed that subjects are unified over time and that this unity stands in need of explanation.

least, some basic persistence over time. As such, there is an inclination to align subjective self-awareness with the minimal sense of subjecthood discussed above, while objective self-awareness is aligned with a richer conception of self as a particular and concrete entity that is unified over time. However, while the division between a thinner and a thicker sense of subjecthood along these lines is appropriate, and endorsed in one way or another by both Kant and Zahavi, it is less than clear that the basis of a subject's persistence over time requires a form of objective self-awareness. In other words, while objective self-awareness might well involve an awareness of oneself as persisting over time, it may also be the case that subjective self-awareness itself already involves such a persistence. Indeed, both Kant and Zahavi seem to also endorse this latter view.

As has already been argued above, subjective self-awareness is at least a necessary condition of any form of subjecthood. Thus, to be a unified self requires that one is, minimally, aware of oneself as the pre-reflective subject of one's experiences. Unless one was presented as such to oneself in experience, beyond the fact that it might not be possible to make sense of the very idea of experience in such a circumstance, there would be no possibility of constituting a unified self over time. Nevertheless, there are at least two potential arguments in favour of the idea that not only is subjective self-awareness necessary for such unity, but that it is already sufficient for such unity; that subjective self-awareness is not just constitutive of subjecthood but constitutive of unified subjecthood over time.

A version of the first of these two arguments is presented by that aspect of Kant's thoughts regarding self-awareness that was not considered above and is, accordingly, transcendental in nature. For, Kant argues, the sense of self as transcendental subject that is achieved through pure apperception is necessarily transcendently *unified*. This is to say, that pure apperception is synonymous with the transcendental unity of apperception. The most important reason being that experience is never simple or of a singular representation, but itself always unified in the form of general experience:

‘There is only one experience, in which all perceptions are represented in thoroughgoing and lawlike connection, just as there is only one space and time, in which all forms of appearance and all relation of being or non-being take place. If one speaks of different experiences, there are only so many perceptions insofar as they belong to one and the same universal experience. The thoroughgoing and synthetic unity of perceptions is precisely what constitutes the form of experience, and it is nothing other than the synthetic unity of the appearances in accordance with concepts.’ (A110)

However, Kant argues, we cannot explain the unity of this ‘universal experience’ unless we presuppose the existence of a unified subject of that experience. This is to say, if to experience is to experience a unity of diverse representations, each of which presupposes as their vehicle an implicit subject, then a unified subject is the transcendental condition of possibility for the unity of

this experience. In other words, if pure apperception itself was not unified we would have a unity of experience and a disunity of experiencing subjects (a prospect which is either contradictory or absurd):

‘The thought that these representations given in intuition all together belong to me means, accordingly, the same as that I unite them in a self-consciousness, or at least can unite them therein... [O]nly because I can comprehend their manifold in a consciousness do I call them all together my representations; for otherwise I would have as multicolored, diverse a self as I have representations of which I am conscious. Synthetic unity of the manifold of intuitions, as given a priori, is thus the ground of the identity of apperception itself, which precedes a priori all my determinate thinking.’ (B134)

The thought is that if experience is unified, and if experience involves an awareness of the subject of experience, then it presupposes that said awareness is of the same subject across experience as the principle of its synthesis. In other words, because experience is never simple but always multiple and interwoven, this presupposes that the awareness of self given in that experience is also unified (otherwise distinctions in the awareness of self would dissolve the supposed unity of experience).

It might be thought that this is irrelevant, considering the fact that Kant’s account of subjective self-awareness has already been rejected. However, it is important to realise that this argument can be accepted without simply returning to Kant’s understanding of subjective self-awareness with which we found fault (or at least Brook’s interpretation of this understanding). Here the Zahavian view of things can be accepted, in which subjective self-awareness is given pre-reflectively, as opposed to transcendently, in the nature of experience itself as its mode of presentation. But having accepted this, and having accepted the idea that experience is in some sense unified, it can be argued that while subjective self-awareness is experientially given, the unity of the self given in subjective self-awareness is transcendently presupposed.

The problem with this argument is that it would appear to entail that while we have pre-reflective awareness of ourselves as subject, we do not have pre-reflective awareness of the unity of that subjecthood over time. Rather, as a transcendental awareness, it would seem that to actually experience the unity of our subjecthood we would have to attend to our experience in exactly the same sense that Kant seemed to think was necessary in the case of pure apperception. Thus, the strength of this argument is going to largely depend on the reader’s intuitions. If one is willing to accept that one’s unity is not given in the pre-reflective presentation of one’s experience as for-me, then this transcendental argument may be convincing. Conversely, if one thinks that the appearance of oneself as unified over time must already be a feature of the way in which experience is presented as for-me, then the transcendental argument will appear insufficient.

Nevertheless, just as with the case of self-awareness, phenomenology offers a non-transcendental equivalent to Kant's thought, which allows us to account for the pre-reflective awareness of subjective unity over time. Indeed, in this respect Kant already uncovers the pertinent point, in recognising that we must understand experience as, at least to some extent, synchronically and, more importantly, diachronically unified. For, in order to be conscious of intentional objects that persist over time, which we obviously are, it must be the case that our consciousness presents us with more than awareness of temporally discrete moments. For instance, to borrow an often utilised example, consider the case of listening to a particular melody. It is difficult to make sense of such experiences if we understand them as simply made up of a series of temporally distinct sounds, that is, if what we actually experienced was a series of unrelated notes which, through some conscious reflective act, were unified into a single intentional object, namely, the melody. If nothing else, this is a possibility that is ruled out by the simple fact that each of these notes would in its own right have to be understood as temporally extended to some extent and would, therefore, give rise to its own need for reflective unification. As such, there is a sense in which we are forced to either accept that experience is, to a greater or lesser extent, temporally extended, or face a dilemma equivalent to Zeno's paradox of the arrow (in which movement is rendered impossible by the infinite divisibility of time). Thus, we must understand experience as immediately grasping the melody as a single temporally extended intentional object; the way in which we experience a melody is inseparable from the fact that it is already, pre-reflectively, presented as such a temporal object. Moreover, unless we grasp the melody as a single temporal object prior to reflection, it is hard to see how this unity could be achieved by reflecting on the temporally distinct parts of the melody after the event.

Nor do we require recourse to intentional objects that are quite so explicitly dependent on temporal extension, in order to demonstrate the necessity of this diachronic unity. For example, in observing the lamp that is sitting on my desk, while the lamp does not change in any overt way from one moment to the next, my understanding of it as the same lamp across the time that unfolds in my attending to it should not be understood as the product of some independent reflective act. I do not need to add-up, as it were, all of the discrete temporal moments in which it appears in my perceptual field in order to comprehend the lamp as extended across time. In part because Zeno's paradox of the arrow would once again rear its head if this were the case, but also, and more importantly, because the temporally extended nature of objects is simply intrinsic to the way in which I perceive them, and my perception of the lamp would be radically different if I did not immediately grasp it in this way.

Explaining exactly how experience presents itself as unified in these ways is perhaps the truly troubling question. As Zahavi puts the problem, ‘how can we be conscious of that which is no longer, or not yet, present to our consciousness?’ (2008, 56) Another way this might be put is by saying that there appears to be a tension between our understanding of the subject as *in the present moment* and the idea that experience is temporally extended. In other words, the subject of consciousness must be understood as both in and moving through time. As such, there is always somewhere that the subject is in time, the present, and somewhere that they are not, the past and future. Therefore, the problem of how experience can be temporally extended appears to be the problem of how these places in time where the subject is not, i.e. the past and the future, can still form an essential part of the subject’s experience in the present. Moreover, as Zahavi quite rightly points out (2008, 56), it is insufficient to point to the role of either memory or imagination to fulfil this task. For, if nothing else, these are usually reflective forms of experience and the appearance of the melody as continuous has nothing to do with, for instance, my building up a continually growing memory of the notes just played or imagining what notes might be about to be played.

However, while many compelling solutions to this quandary have been proposed by phenomenologists, these need not be of great concern here. It is important that a solution to this puzzle can be given, but the details are not crucial for present purposes. All that needs to be realised for now is that our experience of the world must be understood as temporally extended. For, in this respect we have access to an argument that Kant does not. Kant suggests that the unity of the subject is transcendently given in experience, because the unity of experience demands a unified subject of that experience. Therefore, actual awareness of one’s own unity over time requires some form of attending to one’s experiences. However, it has been argued that the subject is not just a transcendental condition for the possibility of experience, but immediately given, or constituted by, the pre-reflective subjective character of that experience. It has also now been acknowledged that this experience in which the subject is constituted must be understood as temporally extended. Thus, it is quite straightforwardly the case that the subject of such experience must itself be understood as temporally unified. For, the nature of the subject follows the contours of the experience in which it is constituted and the nature of experience is temporally extended. To put it differently, if subjective self-awareness relates to the first-personal presentation of experience, then this presentation cannot be understood as temporally discrete, because the intentional objects of which it is an awareness are not themselves temporally discrete.

What this entails, in part, is that the unity of the subject is going to be dependent on the extent of the unity of the experience in which it is constituted, and the extent to which experience is an

unbroken *stream* is open to debate. At the very least it would appear to be a stream that will be broken by events such as dreamless sleep or being knocked unconscious. Nevertheless, once again this ought not to concern us a great deal with regard to present purposes. The purpose of the present discussion is not to present an account of the subject or its unity in every important sense, which would undoubtedly involve going well beyond that which is provided by subjective self-awareness alone. Rather, just as we should think of such self-awareness as constitutive of a subject in the most minimal and basic of senses, so we should also think of it, in virtue of the temporal extension of experience, as constitutive of a minimal subjective unity over time; a unity that will have to be bridged or completed by richer notions of self provided in other ways, but which cannot even be conceived of independently of this most fundamental unity.

V: Conclusions

This chapter began with a consideration of Shoemaker's thoughts on the distinction between the use of 'I' as subject and the use of 'I' as object, where he held the former to be involved in self-ascriptions that are immune to error through misidentification. Shoemaker interpreted this as evidence that such self-references rest on a form of self-awareness that is non-perceptual in nature. However, Shoemaker's own characterisation is unable to progress beyond the purely negative.

By contrast, it was then suggested that Kant's thoughts on pure apperception present a more positive characterisation of such subjective self-awareness, by understanding it in transcendental terms. On Kant's view the subject, in its most minimal and non-ascriptive form, is an abstract condition of the possibility of experience in general. However, Kant's position ran into difficulties when it came to the matter of whether, and how, an awareness of such subjecthood might be made manifest in experience. Specifically, he is either forced to deny that such self-awareness is actually ever manifested in experience or he requires recourse to a form of non-representational awareness, for which he lacks the intellectual tools.

It is in this regard that the success of Zahavi's position becomes apparent. By taking the idea of pre-reflective forms of awareness seriously, he is able to cash out the intuitions that lie at the heart of both Shoemaker and Kant's thoughts on the matter. He is able to provide a positive characterisation of subjective self-awareness, in terms of the first-personal presentation of experience, which captures its inherently non-ascriptive nature, while also being able to explain how such awareness is necessarily manifested in all experience. Moreover, Zahavi presents an account of subjective self-awareness that successfully establishes a basic criterion of subjecthood. For, it has been shown that to be a subject, at least in the most minimal sense, is to be subjectively self-aware. Or, to put it

more accurately, minimal subjecthood is constituted by the pre-reflective first-personal presentation of experience, that is, its for-me-ness. Finally, it has also been shown that it is part and parcel of this subjective self-awareness that it is constitutive of a subject that persists over time, at least to some extent. It is not, as it were, a temporally discrete subjecthood.

This is still undoubtedly, and intentionally, a very thin notion of subjecthood. Nonetheless, it must be recognised that it is sufficient already to present a significant challenge to any potential account of plural subjecthood. For, it would seem that on this account of subjecthood a defence of plural subjects will only be successful if it can be demonstrated that it is possible for groups to be subjectively self-aware, that is, given pre-reflectively as subject in the mode of presentation of their (ontologically collective) experience. In other words, a plural subjecthood would require an account of plural subjective self-awareness.

The highly controversial nature of such an idea should hardly need to be stated. As has already been suggested, when academics began to start taking ontologically collective forms of thought and action seriously a few decades ago, it was precisely ideas such as this one, associated with outdated and extreme forms of collectivism, that they were seemingly trying to avoid. In fact, it would not be unreasonable to characterise almost all forms of contemporary anti-individualism as trying to defend ontologically collective phenomena while denying that groups can be phenomenally self-aware subjects in their own right. But the idea of plural subjective self-awareness, if defended, would appear to entail precisely this possibility. Nevertheless, the aim of the remaining chapters is to defend this possibility; that groups of individuals can, at least potentially, achieve a form of plural subjective self-awareness.

Chapter 4

Four Accounts of Plural Subjecthood

In the previous chapter it was argued that possessing subjective self-awareness is constitutive of being a subject in at least a minimal sense. In the sense in question, the primary condition of being a subject at all is being given to oneself as such in experience. It was also suggested that subjective self-awareness is an awareness of oneself as a temporal entity and, therefore, is constitutive not just of being a subject but of the unity of that subject over time (at least to some extent). As such, with regards to the possibility of ontologically collective phenomena, the question posed is whether or not such subjective self-awareness could be possessed by groups? If subjective self-awareness is constitutive of subjecthood, then it is natural to think that plural subjecthood would have to be constituted by a form of plural subjective self-awareness.

As noted at the end of the last chapter, such an idea seems to be highly controversial for some fairly obvious reasons, in that it involves ascribing to groups precisely those features of consciousness that most believe can only be sustained at the individual level. However, it may not be immediately accepted that the issue ought to be posed in such absolute terms. Perhaps, rather, there are alternative ways in which to defend the possibility of plural subjects that do not involve ascribing phenomenal self-awareness of any form to groups.

With this in mind, four distinct attempts to defend the possibility of plural subjects will be considered below, the first three of which attempt to do so without any suggestion that groups can be subjectively self-aware. In other words, they deny that the possibility of plural subjects rests on the possibility of plural subjective self-awareness. The first two, put forward by Michael Bacharach (section I) and Carol Rovane (section II), both offer arguments in favour of plural subjecthood that would rest on a fundamental re-evaluation of the view of individual subjecthood that was advanced above. As such, Bacharach argues that individual subjects are, contrary to intuition, already a form of plural subject, while Rovane argues that we are wrong to conceive of subjecthood in terms of phenomenal features at all. The third, on the other hand, offered by Margaret Gilbert (section III), is compatible with accepting the underlying view of subjecthood put forward by Zahavi, but argues that groups can count as subjects in a quite different sense. Hence, for Gilbert, there are such things as plural subjects to which certain intentional states can be ascribed, but they should not be thought of as subjects in quite the same sense as individuals.

Nevertheless, it will be argued that each of these accounts fails, and does so as a direct result of their inability or unwillingness to take subjective self-awareness into account. In other words, it will turn out that there are problems with trying to defend the possibility of plural subjects without plural subjective self-awareness. While these particular criticisms do not amount to a conclusive refutation of this line of thought, they will be taken to be at least sufficient reason to take the possibility of plural subjective self-awareness seriously, despite its controversy.

Consequently, a fourth account will be considered, in which Hans Bernhard Schmid (section IV) has directly argued for the possibility of plural subjective self-awareness. Schmid appears to accept both the underlying view of subjecthood in terms of subjective self-awareness and the possibility of such a form of self-awareness being possessed by groups. However, while Schmid's position has notable advantages over the other three and demonstrates the credibility of the idea in general, it will be argued that his account is also not without its problems. In fact, it will be suggested that there is a fault in Schmid's understanding of subjective self-awareness that not only presents a problem for his exposition of plural subjective self-awareness, but actually exposes a quite peculiar flaw with the idea that subjecthood in general is constituted by subjective self-awareness. Specifically, the error that Schmid makes gives rise to a question of individuation, in that it is not clear what feature of subjective self-awareness allows it to constitute a particular individual subject. It will be suggested that this problem forces us to turn back to the idea of subjective self-awareness in order to make it robust enough not only to deal with this problem, but to do so while leaving open the possibility that it might not be stuck at the individual level (i.e. that it might still be open to pluralisation in something like the sense that Schmid envisions).

I: Bacharach

The first account for consideration is rarely, if ever, discussed as a form of plural subject theory, even by Bacharach himself. Moreover, if he does present such an account, it is a position that is wrapped up with his theory of the individual as a team which he began to develop only towards the end of his life. As such, his thoughts on the idea only appear in his final, posthumous, book 'Beyond Individual Choice' (2006), which was pieced together from his notes and early drafts by Natalie Gold and Robert Sugden. What we have on this idea, therefore, is a couple of brief passages in which he presents a rough sketch of the theory. As such, a certain amount of reconstruction is inevitably involved.

It is also important to recognise that Bacharach was not primarily a philosopher, but an economist and game theorist. His main concern in 'Beyond Individual Choice', and his entry into this

conversation, was to defend the possibility of what he called *team reasoning*, primarily as a response to certain fundamental problems in game theory such as the coordination problems that arise, particularly, in the context of games such as *The Prisoners' Dilemma* and *Hi-Lo*. The broad idea of this response is the suggestion that we ought to reject one of the core assumptions of traditional game theory that designates the individual as the sovereign seat of reason and action (i.e. individualism). The core thought of team reasoning, is that there are situations in which a group of individuals, when confronted with a decision problem such as a Prisoners' Dilemma, can pose the problem not for themselves as individuals, but for the group as a whole; they can ask themselves not 'what should *I* do?' but 'what should *we* do?'

In turn, Bacharach's theory of team reasoning can only be made sense of in the wider context of his theory of *framing*, or *variable frame theory*. A view which, in short, attempts to reject the idea that there is any single objectively valid way in which an individual might confront a situation, but rather that the perception and structure of the options available to them will be, at least in part, dependent upon the descriptions and evaluative scales that they employ. It is by extending this theory that Bacharach is able to present his own unique account of team reasoning.

It is fair to say that for Bacharach there are three different *modes* or *levels* of framing. The first is the most basic and foundational. It is what one might call *object framing*³⁹, in which the individual's choice between the various options available to them when confronting a given decision problem (the objects from which they must choose) is going to be partly dependent upon how said individual perceives and describes them. A great many game theoretical experiments are based around trying to present observer or context independent choices, such that the results are easily quantifiable and, therefore, open to measurement and comparison (for instance, money or tokens are often utilised as they possess an easily quantifiable value). *Object framing* denies the very possibility of such choices. Rather, we should not describe the individual as deciding between objectively defined choices at all, but between *act-descriptions* of those choices. Consider the following example described by Gold and Sugden:

Suppose you are taking part in an experiment; you are presented with [the figure below] and asked to choose one of the three symbols by drawing a circle around it.

O Δ X

³⁹ The terminology used to describe the varieties of framing is entirely my own and not utilised at all by Bacharach. It is also important to note that the following is a very brief account that purposely avoids addressing a number of the particularities and nuances of the theory. The aim is only to provide the broad strokes necessary to comprehend Bacharach's resulting theory of team reasoning.

Objectively, there are three distinct symbols on the page, and you must choose one of them. If we are the designers of the experiment, we might identify these three symbols to ourselves as x_1 , x_2 and x_3 (working from left to right), without supposing that, for example ‘ x_1 ’ is a term that has ever occurred to you. Then $\{x_1, x_2, x_3\}$ is a set of objects. But, for you, these objects are described by predicates such as *circle*, *triangle* and *cross*. If you consciously choose to circle the object x_2 , you do so by thinking of this action *as* something for which you have a description—say, by thinking of it as ‘choosing a triangle’. The decision problem that you face can be defined by a set of act-descriptions such as $\{\textit{choose the circle}, \textit{choose the triangle}, \textit{choose the cross}\}$. (Gold and Sugden 2006, 10–11)

In other words, how to interpret the choices that an individual makes when confronted with a decision problem is going to require first and foremost that we understand how the individual understands and describes, or *frames*, those choices. In turn, this is going to largely depend on the evaluative and conceptual tools that are uppermost in the individual’s mind when they address the situation (Bacharach 2006, 69). For instance, while one might consider the above decision problem in terms of the act-descriptions $\{\textit{choose the circle}, \textit{choose the triangle}, \textit{choose the cross}\}$, one might equally do so in terms of $\{\textit{choose the shape on the left}, \textit{choose the shape in the centre}, \textit{choose the shape on the right}\}$, and one might conceivably harbour different evaluative scales when ordering types of shape to ordering placement of shapes⁴⁰.

For Bacharach, however, this is only the most basic role that framing plays in guiding our response to decision problems. The second mode of framing that Bacharach discusses, which is far more important with respect to his overall theory, is what might be called *player framing*. It is the sense in which, for the agent, not only are their options framed in terms of their act-descriptions, but also all of the other players in a game (presuming that it is a game with more than one player), including itself, are also framed. An agent never confronts a given situation in terms of a raw, blank individuality, as some traditional accounts in choice theory would have it, but in terms of the self-conception that they happen to be employing at the pertinent moment (a self-conception that is formed in light of various identities, such as our relationships with other people and our group identifications). Moreover, as Bacharach says, ‘Nor is it the case that a given person has a single self-concept; self-conception is self-framing, and shares the characteristic instability and context-dependence of all framing.’ (Bacharach 2006, 74) We employ a variety of self-conceptions at a variety of times, and understanding how an individual can be expected to act in relation to a decision problem will often be a matter of understanding which self-concept they happen to be utilising at any particular time. Thus, it is the sense of framing in which I never confront a situation simply as

⁴⁰ The example is of course rather sanitised, but it is easy to extrapolate to more concrete examples. Moreover, it is likely that the range of possible sets of act-descriptions is going to be much larger with any kind of *real-world* decision problem.

an *individual*, but, for instance, as a *teacher*, as a *police officer*, as a *parent*, or as a *friend of John*. It is not just the sense in which, for instance, a police officer might have other factors to take into consideration when making choices (such as the risk of losing their job), but the sense in which their very identity as a police officer changes how they conceive of the situation⁴¹.

Just as there is no objective way in which to interpret the choices available to the player in any game, but must be understood in terms of variable act-descriptions of those choices, the set of act-descriptions that is utilised by the player is going to be largely dependent on which self-conception they utilise in relation to the situation. A police officer who, in the course of their patrol, comes across a bank robbery, is likely to view their available options in wholly different terms than if they happened across the same situation while shopping as a civilian on their day off (for example, the former might confront a thief as a target, while the latter might see them as a threat). Likewise, if I am addressing a decision problem in which my co-players are police officers, determining my viable options may be, in part, a matter of determining whether they are more likely to act in accordance with their identifications with the police force or otherwise. As Bacharach wittily puts it:

‘If I see things as a department man at one time and as a family man at another, then at different times I have different goals or, at least, different occurrent goals. There are moments at which my colleagues and I, the whole ragbag of us, see the world through departmental eyes, and moments at which the rags occupy other bags.’ (2006, 75)

While this sense of framing puts paid to the idea of raw, blank individual actors, it is important to note that the argument has still not progressed beyond the level of individuality. Individuals might derive their goals and evaluative standards from various group identifications, but the goals in question are still their own individually sought goals. While Bacharach has provided reasons to challenge the orthodoxy of objective preference formation and choice in game theory, he has not yet provided reasons to challenge the assumption of individual subjecthood enshrined therein.

However, this is where Bacharach’s most radical point is made. He poses the question that, if we accept the point that not only are the available options, but also the players themselves, subject to a process of framing, then why is it necessary that the resulting frame (that is the self-conception) be an individualistic one? Consequently, Bacharach argues for the possibility of particularly strong forms of group identification and the reasoning processes that they lead to. In short, he argues that in some cases what it means to group identify is to view a given problem not as a problem for *me*, but as a problem for *us*. In other words, to utilise a collective, rather than individualistic, frame

⁴¹ This is not to assume that someone automatically has such a strong identity purely in virtue of being a police officer. It is perfectly plausible that someone might be employed as a police officer and yet have no interest in serving the goals of the police force whatsoever.

when confronting a given situation or decision problem. This is to say, if all reasoning occurs from some perspective, why not from the perspective of the group? Strong group identifiers⁴² are led to pose decision problems as problems for the group and, therefore, reason on the basis of what the group as a whole should do and act accordingly. This is what you might call *collective player framing* and it is, for Bacharach, the basis of team reasoning. To team reason is to consider the group that one is part of to be a hypothetical rational actor, to determine what the rational thing for that actor would be to do, and then to perform one's own part in that larger group action.

So far this argument does not go beyond the logical possibility of such collective player framing. It is one thing to say that it is conceivable that groups do often team reason and act accordingly, but it is quite another to show that they actually do. Moreover, a large question remains regarding the relation between the individual and the group in team reasoning. In the case of individual reasoning it seems unproblematic to suggest that the individual is somehow committed to the decision it makes as the result of reasoning. In the group case, Bacharach's theory would seem to demand that somehow the relevant individual members of the group are committed to the decisions reached as a result of team reasoning. But the question is how or why? In other words, it might be possible for all of the individual members of a group to work out what *we* should do, but why should any of them, when considering the matter from their own individual perspectives, care about decisions reached through team reasoning? To put it differently, what is there to make team reasoning efficacious at the individual level?

To this end Bacharach first points to the wealth of empirical evidence that might be taken to support his view of group identification. As discussed, for Bacharach, group identities are the basis for team-reasoning because they dispose the agent to think in collective terms and to pose problems for the group instead of for themselves as individuals. Thus, he takes the fact that there is a great deal of empirical work that supports the idea that group identification is both a particularly strong phenomenon and that it encourages collective action, as *prima facie* evidence in support of his theory of team reasoning. In other words, he argues that his position provides the best explanation of this available evidence.

⁴² Of course, there is a sense in which the police officer considered earlier could be said to group identify when confronted with the bank robbery, in that they are forming their decisions, at least in part, in light of the group from which they derive their particular self-conception (namely, the police force). But in this sense of group identification the self-conception is still individual even if it is formed in light of the group. In the stronger sense being utilised here, however, the self-conception is not individual at all. With respect to the police officer, it would not be a question of considering what one should do as an individual member of the police force, but rather considering what the police force as a collective ought to do.

Nevertheless, and he seems to be well aware of this fact, there are of course alternative explanations of this evidence on offer. While the available evidence certainly does not tell against his account, it does not conclusively tell in its favour either. As such, he then turns to the theory of the individual as a team. The core purpose of this argument is quite clearly to turn our intuitions on their head. Rather than continue to argue in favour of the idea that groups can team reason, he turns the tables and suggests that this should not strike us as strange. For, he proposes, *individual reasoning is already a form of team reasoning*.

He recognises that the main reason the reader might be sceptical of his account of reasoning from the first person plural is because all of our intuitions are geared towards the first person singular. We take our own subjecthood, and the perspective that it affords us, as an *a priori* given and as such are immediately inclined to think that any account of the plural subject or perspective is going to have to meet similar conditions or live up to the same standards. Bacharach attempts to overturn this thought by suggesting that, once we examine individual subjecthood and reasoning in more detail, we discover that the individual is itself already much like a special kind of team and, consequently, that team-reasoning is in fact the norm and individual reasoning is simply a special case of that norm⁴³.

Bacharach argues that, if we consider an individual subject, we can think of them in terms of a series of temporally disparate *time-slices* in virtue of the fact that every subject is temporally extended. Each *slice* has the relevant properties of a subject and is able to reason as to what it ought to do now and in the future (Bacharach 2006, 89). In other words, each time-slice, and consequently each reasoning process, also involves reasoning over the actions of each of the other slices that are relevant to the carrying out of the action (as all action takes some amount of time). As such, and as a direct consequence of our temporal existence, reasoning as an individual necessarily involves reasoning on behalf of, and from the perspective of, a series of disparate individuals (i.e. other temporal parts of myself to which I don't have immediate first-personal access, but which are necessary to carry out the action that is decided upon). As an individual, when deciding upon the best course of action I am not merely concerned to discover what the rational thing for me to do is now, at this very moment of the decision making, but what is the rational thing for me to do over a certain course of time⁴⁴.

⁴³ He even suggests at one point that we might view individual personhood as a special case of self-framing, that is variable depending on the particular past, present, and future moments of ourselves that we choose to include in the frame (2006, 70).

⁴⁴ There are clear similarities between this idea and Bratman's planning theory of agency, discussed in Chapter 2.

Thus, in individual reasoning I do not simply adopt the perspective of myself *here and now*, but that of all my *selves* that are pertinent to the action. The only difference between the case of an individual and the case of a group or team is that the latter is spatially disparate while the former is temporally disparate. However, both involve a plural perspective of some sort and both involve reasoning as to what is rational for the whole, not just the part. In the end, therefore, Bacharach's thought appears to be that we ought not to be too distressed by the problems posed by team reasoning, for any problem involved in team reasoning is a problem for individual reasoning as well, in that the latter is only a variation of the former.

In many respects this is an ingenious solution to the problem of trying to argue that team reasoning in fact occurs. Where others who have wished to argue for ontologically collective forms of thought or action are left with the problem of arguing for some group level ability that looks quite radical, Bacharach takes the bold step of arguing that the ability that he has in mind, that of team reasoning, is in fact characteristic of, or defines, reasoning as a whole; to deny that we do actually team reason would be to deny that even individuals reason.

Nevertheless, it is the very aspects in which Bacharach might be correct that pose the biggest problems for his position. The individual might well be taken as a particular kind of team of temporally differentiated selves, but herein lies the problem. While this might be the case when we take a step back and consider individual subjecthood from afar, it is crucially not the case that the individual considers *itself* to be a somewhat disparate team. Rather, from the first person perspective one's own subjecthood appears as a unified and inseparable whole, at least to a certain extent. When one refers to oneself it is precisely assumed that the *self* in question is not changing from one moment to the next. Of course, it must be recognised that we all change as people on a daily basis and in a wide variety of ways, but the very possibility of such change rests on the supposition that something remains constant; that it is the same being, in some sense, that changes (i.e. that at time t_1 I am numerically identical to, although qualitatively distinct from, myself at time t_2).

This is a problem that is not entirely lost on Bacharach, who in response immediately attempts to turn our intuitions on their head once again by appealing to the notion that both individual and group identification involve a process of *entification*:

'We can partially make good this deficiency... by thinking of subjective personal identity as a special case of group identity in which each temporal agent in the sequence entifies the whole sequence, and conceives itself as a member of this entity. It has thoughts like 'I am Zelda'. This 'am' raises a puzzle. How can it be true that something (a subperson) that *is* part of a person *is* a person? The answer is, I think, that the theoretical entity picked out by the

first 'is' (the agent) is engaging in the talk characteristic of group identifiers. Thinking that she is Zelda is analogous to thinking that we are Zanzibar.' (Bacharach 2006, 89)

Thus, for Bacharach, in the same way that individual reasoning might be considered a special form of team reasoning, personal identity and individual subjecthood might be considered a special form of group identity and plural subjecthood. To personally identify is, as it were, to mentally unify a team of temporally disparate selves in the same way that one might mentally unify the members of a group into a single entity. Individuality is thus the result of a process of group identification with our other temporal moments.

This is the point at which Bacharach is most clearly putting forward what looks, to all intents and purposes, to be a defence of plural subjects. But rather than taking the individual as sacred and trying to extend its subjecthood to groups, Bacharach is arguing that individual subjecthood is already a form of plural subjecthood; that there are, as it were, *only* plural subjects. If individuals are temporally disparate entities, then any sense of self attained by an individual is going to involve some process of group identification (or, to use Bacharach's terminology, entification).

This is undoubtedly an appealing thought and a truly unique way to consider the nature of individual subjecthood. However, as the discussion of the previous chapter demonstrates, that individual subjecthood is something that must be achieved in spite of the temporally disparate nature of human beings, is not something lost on philosophers. If there is such a thing as an individual subject, then, as Bacharach recognises, this subject will involve a series of temporally disparate parts (at least from a particularly detached third person perspective). And, in virtue of this, a persistent sense of self must be understood as in part a construction, in that it will require a process of entification whereby these temporal parts are understood and identified with the same whole. Nevertheless, the truly pertinent question is how or why this process occurs in the case of individuals?

Now, we have already seen one answer to this question, in the form of subjective self-awareness. The argument being that the basis of subjecthood, and the unity thereof over time, is being given to oneself as such in experience. Of course, Bacharach may not agree with this assessment, but he does not provide an alternative. As such, it is not so much that Bacharach's position is either wrong or bad, it is just that it is lacking a crucial argument. In lieu of an alternative account of subjecthood, Bacharach's argument does not present us with a complete defence of plural subjecthood, but merely restates the question with which this chapter began. In other words, if Bacharach wants to argue that individual reasoning and subjecthood are a lot like team reasoning and plural subjecthood, then he needs an argument to demonstrate how subjective self-awareness could be possessed by

groups, for, as is argued above, it is such awareness that is the principle of unification in the individual case.

II: Rovane

It is this last point that is rejected by Carol Rovane. Like Bacharach, Rovane's defence of the possibility of plural subjects is rooted in a re-evaluation of individual subjecthood⁴⁵, but not in the sense that individual subjects are already a form of plural subject. Rather, she wants to provide a particular conception of what it means to be an individual subject and then argue that subjecthood understood in this sense can be naturally applied to groups of humans (as well as, in a sense to be explained, parts of humans). It is, she argues, a faulty conception of subjecthood, one too wedded to the experiential point of view, which prevents us from considering groups to be viable candidates for qualifying as subjects in their own right. As such, unlike Bacharach, Rovane is not lacking an argument for the unity of the self, but rejects the idea that any phenomenal feature is the condition of such unity.

To this end, Rovane begins with an interpretation of Locke as arguing that personal identity is secured by the unity of an experiential point of view⁴⁶; a view with which the argument of the previous chapter has clear affinities. Locke used this definition to argue that personal identity should not be equated with either the identity of a single animal or enduring soul. Rather, our identity is secured by that to which we have direct mental access and, while ordinarily we consider such a point of view to adhere in a particular animal or body, there is no reason to think that this must be the case. It is, Locke argues, perfectly conceivable that our consciousness could be transferred to a different body.

Consequently, Rovane takes Locke to affirm roughly the following three points (1998, 14):

- (1) A person is something with a first person point of view
- (2) The identity of a person consists in the unity and continuity of such a first person point of view
- (3) The first person point of view of an individual person need not coincide with an individual soul or an individual animal.

⁴⁵ Rovane, in taking Locke as her critical starting point, talks almost exclusively in terms of persons and personhood. However, unless stipulated otherwise, I will take 'person' in this context to be synonymous with 'subject' or 'self'.

⁴⁶ Rovane recognises (1998, 14) that this might be a slightly controversial reading, as discussion of Locke's views on personal identity have tended to focus on the role of memory. However, she suggests that this is a slight mistake and memory ought to be taken only as one part of a theory that is committed to identity being equated with a stream of consciousness more generally.

Rovane does not wish to reject any of these points, but believes that they should be defended not by thinking of the point of view in question as an experiential one, as Locke does, but as a rational one. Where a rational point of view is, 'the point of view from which a person deliberates' and where, 'The proper goal of deliberation is to arrive at, and also to act upon, all-things-considered judgements about what it would be best to think and do in the light of everything in the deliberator's rational point of view.' (Ibid. 1998, 21) Presumably, although Rovane rarely makes this explicit, such deliberation is made on the basis of the goals or aims of the deliberator. So, there must be some sense in which a rational point of view is understood in terms of those states of affairs that the deliberator aims to bring about, as that which distinguishes rational agents from one another (a significant point which will be returned to shortly).

On this basis, Rovane then offers what she calls the *normative analysis* of personal identity, as opposed to the predominantly epistemological approach that motivates Locke and has its roots in Descartes (Ibid. 1998, 210). Rovane does not deny that there are unified experiential points of view in the sense that Locke envisioned. It is just that such a point of view should not be equated with personal identity. The primary reason for this is that Rovane takes 'person' to be an ethically loaded term, in that the ascription or denial of personhood is one of the most ethically significant acts. As such, she suggests that ethical considerations ought to play a large part in how we define personhood, which leads her to put forward the *ethical criterion*. This criterion stipulates that:

'[P]ersons are agents who can engage in agency-regarding relations. That is to say... that persons can engage in the particular kind of relation which arises between agents when one agent attempts to influence another, and yet aims not to hinder its agency.' (1998, 72)

Rovane takes it as no simple matter to define what it means to take part in such *agency-regarding relations*. Indeed, a more thorough consideration of Rovane's position would require an in-depth analysis of this definition, on which Rovane herself clearly places a great deal of emphasis. For present purposes, however, it is sufficient to surmise that the pertinent point is that Rovane takes agency-regarding relations to relate to the way in which persons are able to rationally influence one another. This is to say, to take part in agency-regarding relations is to be capable of recognising what will act as a reason for a particular person and, thereby, being able to submit them to rational influence via an appeal to this reason. But in order to recognise what will act as a reason for that person, the influencer must attribute a certain rational point of view to them on the basis of which such a reason can be deduced. The influencer must put itself in the rational shoes of the one to be influenced, as it were. In turn, Rovane argues that this ability requires that the influencer also has a rational point of view from which to do the influencing; if the influencer does not first have a rational point itself, then the process of attributing such a point of view to another, and then

adopting this point of view for the purposes of working out what will act as a reason for them, cannot even get off the ground. In a sense, therefore, Rovane takes agency-regarding relations to require a sort of meeting of rational perspectives. Consequently, if the ethical criterion defines a person as one who can engage in agency-regarding relations, and engaging in agency-regarding relations requires a distinct rational point of view, a person is ultimately defined as anything with such a rational point of view. Meeting this criterion is vital for Rovane in the sense that ‘it helps to define and expose a pernicious form of prejudice against persons, which consists in a hypocritical denial of their personhood.’ (1998, 128)

Nevertheless, Rovane recognises that the ethical criterion alone is not sufficient to demonstrate that a rational point of view, rather than an experiential one, is more important when it comes to personal identity. She also needs to show that a rational perspective in this sense is capable of becoming divorced from a particular experiential perspective. The assumption of Locke, and indeed seemingly all those who have defended positions that resemble his, is that an experiential point of view is the precondition of a rational point of view. Rovane takes this view, that, for instance, someone is only committed to consistency amongst their beliefs because these beliefs all feature in one and the same consciousness, as a ‘nearly universal assumption’ (1998, 20). And, indeed, such a view would seem to be entailed by the view of subjecthood extolled in the previous chapter. It is crucial to Rovane’s argument that she rejects this identification because if the rational is held to come with the experiential point of view, then there is no reason to favour the former over the latter when it comes to personal identity. The ethical criterion may define a person as something with a distinct rational point of view, but if having a rational point of view is dependent on first having an experiential one, being a person would still be a matter of experiential continuity, albeit indirectly.

This is where Rovane turns to the suggestion that both groups of humans and parts of humans are capable of having a unified rational point of view in their own right, even though they do not have distinct experiential points of view. Phenomenal awareness, Rovane thinks, is restricted to the level of the individual organism, but rationality is not. Individual humans can be persons, but they can also be the site of multiple persons and form part of a larger person that includes other humans as well. For the sake of clarity, let us call these three possibilities (as Rovane seems to for the most part) *individual persons*, *multiple persons* and *group persons*. While Rovane accepts that the individual person is going to be the paradigm case of having a rationally unified point of view, she argues that both group and multiple persons can still achieve relative degrees of unity and disunity.

If we take the case of group persons to begin with, Rovane holds that their possibility rests on the capacity for *projection*. As was briefly suggested above, the ability to take part in agency-regarding relations in subjecting others to rational influence, rests on the ability to project oneself into the rational point of view of the other (Rovane 1998, 137). This is to say, in order to rationally influence another I need to work out, from their rational point of view, what will influence them. In itself this does not require one to ever abandon one's own rational point of view, but it does require that, 'persons must view their ends and actions *both* from their *own* and others' rational points of view.' (Rovane 1998, 138) However, it is being able to project oneself into others' points of view that enables one to also be capable of seeing things from a shared point of view:

'In joint activities groups of persons deliberate and act together for the sake of common ends. And unlike agency-regarding relations and rational influence, these activities do not require that persons project themselves all the way into another person's own rational point of view so as to take up that person's perspective. These activities require rather that persons project themselves into a rational space that is generated by the ends which they hold in common—a space where distinct rational points of view overlap. When persons project themselves into this common rational space, they can reason and act from the perspective of their common ends.' (Rovane 1998, 138)

In many ways, therefore, Rovane's defence of group persons has many similarities with Bacharach's defence of team reasoning. As was seen, for Bacharach, being able to team reason rests on the ability to consider things from the perspective of the group and ask 'what should we do?' Rovane too is suggesting that humans have the ability to project themselves into certain group perspectives and deliberate from the rational point of view that such a perspective provides. The problem that this posed for Bacharach, however, was to demonstrate the sense in which the individual is committed to reasoning from such a group perspective. Even if one does reason from the perspective of the group, why should that matter to oneself qua individual? It is reasonable to think, therefore, that Rovane's account is going to face a similar problem. After all, she thinks it is possible to project oneself into the rational perspective of another individual person without abandoning one's own rational point of view. Why should it be any different for projecting oneself into the group point of view?

One response is to point out that the rational point of view of the group must be understood as a *shared* point of view; it is a point of view that exists in the overlap between the various group members' individually rational points of view. But it is also important to emphasise that for Rovane, on the definition of personhood supplied by the ethical criterion, there is no priority given to the individual:

‘[H]uman beings don’t actually stand under any rational obligation that would preclude group and multiple agency; and any rational obligation to which human beings are subject is one to which group and multiple agents would be subject as well.’ (Rovane 2004, 190)

A person, individual or otherwise, is just anything with a rational point of view, and in the case of the group, it has a rational point of view because of the goals or ends that the relevant human members share in common. Rovane provides the following example:

‘Imagine that there are two human-size undergraduates who set about to solve a complex philosophical problem together. Imagine further that they together mount a single argument of just the sort that an individual agent might produce on its own... [O]ne undergraduate works out the consequences of what the other has just said at a given stage, and then contributes the next stage in a cooperative constructive conversation... In the context of their joint endeavor, there is an intuitive sense in which it is as if the two undergraduates had only one rational point of view, for by hypothesis, they deliberate and act together just as one agent might act on its own.’ (138-139 1998)

Of course, Rovane is not suggesting that the two undergraduates meet quite the same standards of normative unity that the average individual does. As she puts it (Rovane 1998, 139), they do not need to achieve *overall rational unity* just in order to solve a simple problem, but they do have to achieve *some degree* of rational unity:

‘[I]t is clear that a group of human-size persons, such as the two undergraduates in this example, can achieve a significant degree of rational unity among them. And what is more, the members of such a group can have good reason to achieve such unity together, insofar as they have a shared practical commitment to a joint endeavor whose success depends on it.’ (Rovane 1998, 140)

Moreover, while the undergraduates do not need to attain a state of complete rational harmony, it is perfectly conceivable that we could imagine an arrangement in which a group of humans was committed to achieving overall rational unity such that they did operate as a single individual. Either way, what is crucial is that a person, for Rovane, is identified with a rational point of view and groups of humans, despite what Rovane takes to be their clear experiential discontinuity, are capable of possessing such a rationally unified point of view to a greater or lesser extent.

The other way in which this can happen, for Rovane, runs in the opposite direction; we can have several distinct rational points of view within the same human. The most extreme way in which this can happen is demonstrated by *Dissociative Identity Disorder* (DID; a condition that used to be referred to as *Multiple Personality Disorder*). Rovane does, quite rightly, preface the discussion of this disorder by pointing out that there is very little about its analysis that is agreed upon in a clinical sense, and what precisely is going on in such circumstances is open to numerous interpretations. However, Rovane suggests that it is at least possible to make the case that the various personalities exhibited

by someone suffering from this disorder can, in fact, constitute distinct persons or selves in their own right, if we understand these in terms of rational points of view:

‘It is possible and, indeed, it is effective to address and engage such alter personalities separately. When we do so, we attribute to them beliefs, values and choices. And we may subject them to the same principles of rationality that we apply to human beings. We may, for example, point out their inconsistencies and non-sequiturs to them, fully and legitimately expecting appropriate uptake and self-critical response. Whenever we do this, we are recognizing the possibility of alters as agents in the only sense that matters—and human size is once again a traditional and tempting distraction to be resisted.’ (Rovane 2004, 188)

So, there appear to be good reasons to think that such alter personalities operate from distinct rational points of view. However, they are also, Rovane suggests, quite often experientially continuous:

‘[S]ome alters have phenomenological access to the thoughts of other alters—which is to say, they have phenomenological access to thoughts that they do not regard as their own. Very often, there is an asymmetrical relation between the personality who is regarded as the first and dominant one and the alter personalities who appear later; for the latter often have direct phenomenological access to the thoughts of the former (though not always to the thoughts of one another), while the former remains ignorant of the others’ existence.’ (Rovane 1998, 171)

If correct, once again we are faced with a case where a rational point of view is divorced from phenomenal awareness: one experiential point of view with multiple rational points of view. Nevertheless, Rovane is keen to stress that things do not have to be quite so extreme in order for the case of multiple persons to arise. As was suggested in the case of group persons, the commitment to a rationally unified point of view admits of degrees and there is no reason to think that this will be any different when it comes to multiple persons. What is important about cases of DID is that each personality is a separate person just in case there are distinct rational points of view in play. And it is not difficult to imagine a less extreme case in which a single human can embody relatively distinct rational points of view at different times. For Rovane, humans regularly compartmentalise their pursuits in life (Rovane 2004, 190) and the extent to which they do so is the extent to which they constitute distinct persons with distinct rational points of view.

A particularly interesting example of this that Rovane considers (2004, 193–194) is the case of *Donnie Brasco*. The film by the same name is based on an actual case of an undercover FBI agent named Joseph Pistone, who infiltrated the mafia in the mid-70s under the alias Donnie Brasco. As Rovane notes, Pistone’s undercover work forces him to lead a double life that becomes increasingly pronounced as the movie develops. As things progress it becomes clear that Pistone identifies with many aspects of mafia life and develops some deep friendships therein. As such, he increasingly becomes personally involved with the mafia and more distant from his civilian life and family.

Rovane's claim hinges on the fact that for a significant part of the movie Pistone seems to give up on the idea of leading a single unified life. Instead he pursues his life as a mobster and as a family man relatively independently of one another. This is to say, that when he is going about his life as a mobster he needs to pay little or any thought to the fact that he is also an FBI agent or family man, and vice versa. Consequently, it would make perfect sense to engage with Brasco and Pistone as relatively autonomous rational entities. Of course, this would not be possible to quite the same extent as with the alter personalities of someone suffering from DID. After all, Brasco obviously does not completely forget that he is also an FBI agent while fulfilling his role as a mobster, otherwise his value as an undercover agent would be entirely redundant. However, for the most part it would be possible to engage with Brasco as a mobster without ever having to take into account his personal life. As such, once again we have a case where the rational does not go hand in hand with the experiential, as we are faced with rationally distinct points of view within a single human with a unified experiential point of view (as Pistone and Brasco obviously share a single stream of consciousness).

By defending a conception of personal identity as being constituted by a rational point of view, therefore, it appears that Rovane's account allows us to attribute subjecthood to groups (and parts of humans) quite naturally. On this view, subjects come in a range of shapes and sizes, as it were, and plural subjects are just one such form. However, this is only the case if it is accepted that her examples illustrate cases where the rational is capable of becoming divorced from the experiential and, as will shortly be argued, there are significant reasons in each case to deny that this is what they illustrate.

What Rovane fails to recognise is that any point of view, rational or otherwise, can only be understood as arising on the basis of subjective self-awareness. Rovane fails to recognise this because, presumably following Locke, she does not distinguish a minimal sense of self from a thicker or more concrete sense. As outlined in the previous chapter, such a distinction is a crucial part of Zahavi's defence of the idea that subjective self-awareness is constitutive of subjecthood. The point being, that there is a thick sense of self as a concrete entity to which particular properties can be ascribed, but there must also be a thinner sense of self in play, on which this more concrete subjecthood can build; a sense of self as merely something to which properties *can* be ascribed; an awareness of oneself as a pure or simple subject. Without this more minimal sense of self the question of personal identity cannot even arise. To put it differently, in a world in which experience was not presented as intrinsically first-personal (if such a world is even comprehensible), why would it occur to anyone or anything to ask of themselves whether they are a persisting entity? Why would

the concept of subjecthood arise if experience were not subjectively laden, that is, if it existed for no-one in particular?

The point of making this distinction is not to re-affirm Locke's position regarding experiential continuity, albeit with a thinner understanding of the continuity in question. Subjective self-awareness is obviously a feature of phenomenal awareness, but emphasising its significance is as much a criticism of Locke as it is Rovane. Both of their positions set out to understand personal identity in terms of the continuation of a set of particular properties. For Locke it is a set of intentional episodes to which one has direct mental access. For Rovane, it is a set of intentional episodes circumscribed by a practical commitment to a particular unifying project. Zahavi's point, however, is that it does not matter how one wants to think of personal identity in this thicker sense, because in any case a thicker notion of self must rest on a more minimal sense of oneself as the pre-reflective subject of experience; a sense of self that does not require the ascription of any particular properties whatsoever. The experiential continuity in question, therefore, should not be understood in terms of any set of intentional episodes, for all such episodes share in common the quality that marks them out as constitutive of a subject: for-me-ness.

Rovane does seem to have a potential response to this. She can point out that while this minimal sense of subjecthood outlined by Zahavi is all well and good, it is still not sufficient for a rational point of view, and it is a rational point of view that is significant for full-blown personhood. This is to say, subjective self-awareness might be necessary for the question of personal identity to arise, and might even set the realm of those intentional episodes that can be included in a rational point of view, but such a point of view is still not determined by this underlying self-awareness, in the sense that the latter does not govern what is and isn't to be included in the former. For, while every intentional episode will be presented with an intrinsic for-me-ness, not every intentional episode will or should be included in a particular rational point of view. Rather, as she suggested, what is to be included in a particular rational point of view is settled by the goal that such a point of view is employed in the service of; its *unifying project*:

‘The normative analysis says that an individual person exists just in case there is a set of intentional episodes standing in suitable relations, and the set includes practical commitments to unifying projects that bring in train a commitment to overall rational unity within the set.’
(Rovane 1998, 183)

In other words, as was suggested earlier, what it is rational to do is always relative to a certain goal, so it is the goals of the individual that settles what is to be included within their rational point of view.

However, in light of the concept of subjective self-awareness, this way of determining what is included in any given rational point of view must now be recognised as circular. If, for instance, what determines the boundaries of my own rational point of view are my goals, we must then ask what makes these goals *mine*? It is not clear how this question could be answered without appealing to an underlying sense of subjecthood. And, of course, subjective self-awareness gives us an easy answer to this question: what makes it my goal is that I experience it as mine or for-me.

To make this criticism clearer let us consider how it affects Rovane's account of group and multiple persons. To begin with the latter and the case of Donnie Brasco, Rovane's assessment that Brasco and Pistone are actually two separate persons, to at least some extent, is based on the thought that they have distinct rational points of view. As has just been discussed, the intentional episodes that are to be included in their distinct rational points of view will be determined by the relatively independent goals in question. On the one hand, Brasco's rational point of view is determined by something like his goal of being a good mobster, while Pistone's is determined by his goal of being a good husband and father. However, the question at this point is what makes the respective goals in question the goals of two distinct persons, rather than the same overarching individual? Any appeal to their distinct personhood would of course be circular as this personhood is, for Rovane, determined in part by the goals in question. But it is not clear that Rovane has any alternative answer to provide.

Understanding the situation in terms of subjective self-awareness on the other hand gives us a different view of things, and one which seems far more in tune with the facts of the case. The account of subjects offered by subjective self-awareness suggests that there are only two people if their experience presents them to themselves as such, and it is hard to believe that this is going to be the case in this example. As such, what we have is not two distinct rational points of view and two persons within the same experiential point of view. Rather, what we have is a single subject who happens to be deeply rationally conflicted. On the one hand he does not want to abandon his wife and children, but he also clearly identifies with the code of conduct of the mafia and enjoys the system of respect that comes with it. To put it flippantly, what we have is a single person who wants to have his cake and eat it too. That this is the case is shown by the fact that many of Brasco/Pistone's attitudes in the movie would be unintelligible unless he acknowledged the inherent conflict in his life. For instance, it is clear both that he feels guilty towards his family for abandoning them and that he feels guilty with respect to his closest friend, within the mafia, when his duties as an FBI agent, ultimately, force Brasco to betray him. Of course, Rovane is correct that Pistone has two distinct goals that he can pursue relatively independently of one another, but that they are two

somewhat conflicting goals of the same individual is shown by the attitudes he possesses to the relation between them. If nothing else, it is shown by the fact that once in a while he realises that he has to put his mafia duties aside and go visit his family and that he then realises when he has to leave them and go back undercover. Such recognition is only possible if these two goals are conceived of as two rationally conflicting parts of a single rational point of view.

By contrast, there are good reasons to agree with Rovane when it comes to the assessment of people suffering from DID, in that such cases might demonstrate that it is possible for multiple persons to inhabit a single human body (or at least this is a viable, if radical, interpretation of the disorder). However, the extremity of such a condition allows for a much easier explanation; such persons do not just have distinct rational points of view, but they are given to themselves in self-awareness as distinct subjects. In other words, they are distinct experiential subjects and each personality will have a separate sense of their experience with its own intrinsic for-me-ness. That this is the case is shown by the fact that such alter personalities do not recognise the same sort of linking responsibilities that Joseph Pistone recognises. Pistone knows when he has to switch between his different roles and consciously chooses when and where to do so. There is no such conscious decision making in the case of the switch between alter personalities and, if there were, it would surely represent a distinct form of pathology.

As suggested above, Rovane rules this interpretation out on the basis that alter personalities can be co-conscious. Thus, Rovane's thought, contra Locke, is that such a human has experiential continuity but rational discontinuity. However, Rovane's thoughts on this matter are quite telling, in that her description of being co-conscious is that alter personalities 'have phenomenological access to thoughts that they do not regard as their own.' (1998, 171). In other words, Rovane herself makes clear that such mental access will not appear to the subject in question as first-personal. How we then interpret what is going on when alter personalities are co-conscious in this sense is open to debate. However, the most straightforward option is to simply argue that the various personalities of someone suffering from DID are able to access the thoughts of one another in a way that is not usually available to humans. This is to say, that instead of having to infer the mental states of one another from, for instance, facial expressions, they enjoy direct introspective access to each other's mind. But the end result will be the same as it is for an ordinary person when considering the thoughts of another: a first-personally presented experience of someone else as having a certain thought.

Either way, the important point is that on our account, it is the first-personal givenness of experience that is constitutive of subjecthood, not just the ability to directly access such thoughts.

Locke's definition of personal identity in terms of crude experiential continuity may not be able to capture this, but that is because Locke, like Rovane, does not recognise the more minimal sense of subjecthood that comes with subjective self-awareness. Once we do so, we can see that cases of DID do not just differ in the degree of rational unity from cases such as Donnie Brasco. Rather, they differ at a far more fundamental level, in terms of the subjective character of their experiences.

Finally, let us turn to the case of group persons, which is after all, in enshrining Rovane's particular account of plural subjects, our main concern. The problem that the above criticisms pose for this account should now be relatively clear. For Rovane, the rational point of view of a group and, therefore, its distinct personhood, is determined by the unifying project that the group acts in order to achieve. In other words, their personhood as a group is determined by their joint goal or common ends:

'[Unifying projects] provide the principle by which intentional episodes are distinguished as belonging to a particular group person's rational point of view, for it is in the nature of such a project to dictate how the efforts of a group of human beings must be coordinated so as to carry it out.' (Rovane 1998, 173)

The group is a person because it has a rational point of view and the intentional episodes to be included in this rational point of view are determined by the goal of the group. But what makes such a goal the group's? What makes it *our goal*, as opposed to, say, *my goal and your goal*? The latter on Rovane's account, as a series of individual goals (even if they share the same content), would only determine our own individually rational points of view and, therefore, our own individual personhood. Alternatively, any appeal to our group personhood would obviously be circular as this personhood is in part defined in terms of the goal in question. Once again it seems that Rovane fails to capture the way in which a point of view can be distinguished without taking into account subjective self-awareness and she is forced to presume the subjecthood that her account is meant to explain.

III: Gilbert

On the evidence so far, it would seem that trying to defend the possibility of plural subjects through a re-evaluation of individual subjecthood is problematic. Therefore, if plural subjects are subjects in the same sense as individuals, then their possibility would appear to rest on the possibility of plural subjective self-awareness. However, perhaps plural subjects are not subjects in the same sense as individuals. Margaret Gilbert has defended a conception of plural subjects along these lines; one where the subjecthood of groups is different in kind from the subjecthood of individuals.

In fact, there is a temptation to think that Gilbert is not putting forward an account of plural subjecthood at all, despite describing it as such⁴⁷. For the most part she appears to place very little emphasis on the term and does not seem to view her account as confronting certain fundamental questions of subjecthood (such as those considered above). For Gilbert, as will be spelt out in greater detail shortly, to be a plural subject is for a group of individuals to be *jointly committed to doing something as a body*. There is a great deal to unpack in this statement, but what should be clear is that this bears no resemblance to what it means to be an individual subject. On no account does being a subject in the individual sense involve some sort of binding commitment to oneself. Being a subject might create a commitment of this nature, but it would certainly be strange to think that it could be constitutive of that subjecthood in the first place. Perhaps, therefore, ‘plural subject’ in this context is just being used in a loose sense by Gilbert, and we should not place too much emphasis on the subjecthood of such group agents. However, there are good reasons to reject this interpretation. Whether Gilbert explicitly recognises this or not, her account does raise certain fundamental questions of subjecthood. A fact which, as will be shown, manifests itself in the form of the analogy of doing something *as a body*.

In order to recognise the import of this claim it first needs to be understood what it means, in Gilbert’s view, for a group of individuals to jointly commit to doing something as a body, and how such a commitment brings into being something that can be legitimately referred to as a plural subject. There are two key elements to this position: the idea of being *jointly committed* and that to which the commitment relates, namely, the act of doing something *as a body*.

To begin with a consideration of what it means to be jointly committed. Gilbert’s defence of plural subjects is largely guided by the observation that doing something together, in the strongest sense, intrinsically involves certain normative features. Specifically, we find that in acting with others towards a common goal we each enjoy a certain standing to rebuke and make demands of the others involved in the action (Gilbert 2008, 103–106). As such, her account is built to be normative from the ground up, as it were, in order to present a theory in which such relations of obligation are the direct result. It is these obligations, or a certain normative constraint if you will, that holds the group together as something more than the sum of its parts. Gilbert achieves this fundamental normativity by appealing to the idea of a certain foundational commitment that acts as the starting point for any properly joint activity (that is, an ontologically collective form of action). Each of the relevant individuals begins by expressing a certain conditional readiness to jointly commit to doing

⁴⁷ David Velleman (1997) has, for instance, made the suggestion that perhaps Gilbert is only using the term ‘plural subject’ as a *façon de parler*.

something together (Gilbert 2008, 138–140) and it is known as a matter of common knowledge amongst the members of the group that each has expressed their readiness in this way (2008, 121, 138–139). The nature of this conditionality is that none are committed until all are committed and, as a result, once in force the joint commitment cannot be unilaterally rescinded (1990, 7; 2008, 141–144). This is to say, that while the expressions of conditional willingness are wholly individual, what they represent is a willingness to enter into a commitment that is wholly collective. Once the conditions of the joint commitment are met, that is, that they have all expressed their willingness and this is common knowledge amongst them, they have brought into being a commitment that holds over all of them as individuals, but which does not belong to them as individuals.

Thus, the foundations of plural subjecthood are, for Gilbert, essentially communicative and normative. The normative constraints that provide a standing to rebuke other members of the group when acting together are a result of the fact that as a group they are committed to acting as a group. Just as one might chastise oneself for not abiding by a personal commitment or intention to do something, we can rightfully chastise other parties to a joint act if they do not abide by our commitment to act accordingly (Gilbert 2008, 147–164). Plural subjects are, therefore, founded on acts of communication that commit the individuals, jointly, to doing something together in the sense of acting as if they were a single subject.

It is this notion of being jointly committed that has been the primary source of criticism of Gilbert's position. One such attack, and perhaps the most compelling, has been the charge of infinite regress. Forming a plural subject on Gilbert's account, as has just been noted, requires a joint commitment, but forming a joint commitment, so this line of attack goes, is something that would already require a plural subject. Primarily this argument rests on the claim that communication is already a form of acting together in the relevant sense. As Schmid puts it:

‘If the prospective members of a plural subject mutually have to express their willingness to be jointly committed, in however implicit a way, they have to communicate. Communication, however, is itself a joint action. According to Gilbert's account, joint actions involve a plural subject. Thus, for a plural subject to come into being, another plural subject (the subject of communication) already has to be in place, which sets off an infinite regress’. (2014b, 19)

However, this criticism rests on the claim that communication is a joint action in the ontologically collective sense intended by Gilbert (i.e. not understandable in terms of only individual thoughts or actions), and this is by no means uncontroversial. Of course, the suggestion that communication or language presupposes some form of collaboration or sociality can draw on a rich tradition of thought that, perhaps most famously, includes Wittgenstein and is taken in phenomenology to be

virtually axiomatic. This is also a claim that has been supported by Michael Tomasello's recent work in developmental psychology:

'In beginning to acquire linguistic symbols... infants again demonstrate an understanding of the different but complimentary roles in a social interaction—in this case an interaction involving the exchange of communicative intentions embodied in conventionalized actions.' (Tomasello et al. 2005, 683)

'What could it mean to say that language is responsible for understanding and sharing intentions, when in fact the idea of linguistic communication without these underlying skills is incoherent?' (Ibid 2005, 690)

Nevertheless, this is neither an uncontroversial position nor is it ignored by Gilbert. As already discussed in Chapter 1, she dedicates a fairly lengthy chapter of *On Social Facts* to a thorough discussion of something like this issue (1989, 58–145). Specifically, she takes it as imperative to argue that language, and meaningful behaviour or intentional thought more generally, does not presuppose that the user is socially embedded, in order to defend what she describes as the *intentionalist programme*. This being the essentially Weberian idea that we can explain all social formations in terms of the constitutive role of individual thought, which is prior to these formations. In order to defend this view, specifically directed against the arguments of Wittgenstein and Peter Winch, she draws on an argument made by Kripke and emphatically concludes that:

'[O]ur concept of thought and language in general is such that even an individual forever isolated from others could, in principle, think and use a language. I have argued, if you like, that our basic notion of thought and language in no way involves a community or society of language users.' (1989, 132)

In other words, neither language formation nor language use is necessarily a joint sort of action for Gilbert. It might be countered that while language use *per se* does not necessarily involve such joint action, acts of *communication*, in which one is trying to actively convey a certain meaning to another, might well do. However, once again, this seems contentious and there is no reason to think that Gilbert would be forced to accept that this is the case. While my own intuitions on this matter lie with Schmid rather than Gilbert, it must be taken to be at least problematic to criticise Gilbert on these grounds. It is not as if it is an obviously bad position to hold, and Gilbert is careful to take the time and effort to defend it in great detail (a fact to which Schmid's own criticisms do not do justice).

So, leaving aside the idea of being jointly committed for a moment, let us turn now to the second element of Gilbert's position, that is, the object of joint commitment. Despite the fact that the bulk of analysis, both commentators' and her own, focuses on the act of being jointly committed, the object of such commitments must be taken to be equally significant in terms of the structure of

Gilbert's theory. For, being jointly committed to doing something is not sufficient to constitute a plural subject. To see why this is the case, consider the example of two siblings forced to share a room who, after much squabbling about the violation of personal space and property, decide that it would be better if they divided the room in half and each stuck to their respective side. It is reasonable to imagine this being achieved by each expressing their conditional willingness to stick to one half the room provided that the other also expresses their willingness to do so. Once they have so expressed, on Gilbert's account, they would both be party to a joint commitment to stick to their own side; a commitment that holds over both of them and cannot be rescinded unilaterally. However, it would be strange to describe the two siblings as forming a plural subject just because they are committed in this way. For, what they are not committed to doing in any meaningful sense, is acting *together* and each could happily abide by the joint commitment without ever giving a second thought to the other. As such, joint commitment alone cannot be sufficient for plural subjecthood.

Rather, it would seem that there needs to be some sense of togetherness placed within the object of the joint commitment in order for a plural subject to be formed. This much is shown by the mere fact that Gilbert feels the need to stipulate anything regarding the object of joint commitment. Occasionally she refers to the object in question as a *pooling of wills*, but she recognises that this is a rather vague notion (1990, 8). Far more frequently, and as already suggested above, she refers to the common form of joint commitment as a commitment to doing something *as a body*:

'Joint commitments all have the same general form. People may jointly commit to accepting, as a body, a certain goal. They may jointly commit to intending, as a body, to do such-and-such. They may jointly commit to believing or accepting, as a body that such-and-such. And so on. The general form of a joint commitment then, is this: the parties jointly commit to X *as a body*.' (2008, 137)

In other words, the necessary *togetherness* that needs to be part of the object of a joint commitment, in order to generate a plural subject, is to be found in this idea of doing something as a body. Whatever it is that the group is jointly committed to doing, they must be committed to doing it as a body.

However, in light of what has been said so far on the nature of subjecthood, we are presented with the means to make what amounts to a much more damaging criticism of this aspect of Gilbert's position than the charge of infinite regress considered earlier (one, which, as far as I am aware, largely goes unnoticed). Specifically, the idea of doing something *as a body* is a rather ambiguous analogy. It immediately gives rise to the question of the sense in which 'body' is being used here, and what it means to act as a single such entity. Needless to say, the realm of *the body* is an ontologically contentious area and, at the very least, is going to be open to either a subjective or an

objective reading. In other words, there is a serious question as to whether the body should be understood in this context as a purely physical entity, fundamentally divorced from the realm of the mental and capable of being brought into the purview of subjecthood only through a process of reflection? Or whether this dualistic distinction should be rejected and one's mental life be taken to already be an intrinsically embodied way of experiencing the world? In which case bodiliness must already be understood as part and parcel of what it means to be a subject in the sense provided by subjective self-awareness⁴⁸.

It is not clear what Gilbert's response to such questions would be, but neither option is without its problems. To begin with the latter, more subjective, reading of 'body', it appears to leave Gilbert's position with a rather large theoretical gap. For, if she intends 'body' in a subjective sense her account amounts to the claim that to be a plural subject is for the respective members of the group to be jointly committed to acting as if they were a single unified embodied subject. But it has been argued above that to be a subject, in even the most minimal sense, is to be subjectively self-aware; to be given to oneself as a subject in the mode of presentation of one's experiences. Of course, Gilbert does not want to say that plural subjects are subjects in this strong sense, but then what does it mean to say that the members of the group are committed to acting *as if* they were such a subject without actually being such a subject? To put it differently, what does it mean for a group of individuals to behave as if they were part of a single subjectively self-aware entity? It is hard to believe that convincing answers to these questions could be provided and Gilbert certainly doesn't provide them herself.

However, the objective reading in which 'body' is understood simply as a physical entity (albeit one of a particular form and nature) would appear to fare no better. In contrast to the previous reading this would amount to the claim that to be a plural subject is to be jointly committed to acting as single object or thing. This might immediately seem counter-intuitive considering the fact that Gilbert is trying to defend a theory of plural subjects and, indeed, on a certain *prima facie* reading such an interpretation would appear to lead to a fundamental flaw. The entire point of Gilbert's theory is to provide a defence of ontologically collective forms of thought and action. However, on this reading of 'body' what is to prevent an ontologically distributive analysis? What is it that makes the body in question, as an object, more than the sum of its parts? When we consider the bodies of individuals as vehicles of action we do take them to constitute a sort of natural unity, but this is presumably only because we link such bodies to a single unified subject to which the body shares a

⁴⁸ See Chapter 5 for a detailed response to such questions.

special link or bond. Such a subjective-link, however, is exactly what cannot be assumed if it is our commitment to acting *as a body* that is supposed to establish our plural subjecthood in the first place.

The natural response to this, and this is perhaps the most natural reading of Gilbert's position, is to refer back to the *jointness* of commitment in order to establish the unity of the body in question; that there is a sort of reflexive arrangement in place whereby the joint commitment plays the unifying role with regard to the *collective body*, in the way that subjecthood plays this role in the case of individual bodies. But this presents us with what looks to be a circular argument. Earlier it was suggested that being jointly committed was not sufficient to constitute a plural subject because a group of people could, potentially, be jointly committed to doing something in a wholly non-collective sense, such as the two siblings who commit to dividing their room in half. Rather, a group needs to be jointly committed to doing something in a joint or collective sort of way, which for Gilbert is expressed in the form of them doing something as a body. But now we are asking what makes doing something as a body a joint or collective sort of act and the answer under consideration is that acting as a body is a joint sort of act in virtue of the jointness of the commitment made by those who constitute the body. As such, we appear to have come full-circle. A joint commitment to doing something can be constitutive of a plural subject because its object contains a joint aspect (the idea of doing something together), but this joint aspect is in turn to be explained by the jointness of the joint commitment. In which case, the significance of saying that a group must be committed to doing something *as a body* seems to have been lost entirely.

What this entire problem boils down to is the lack of a principle of unification. Gilbert famously suggests, drawing on the work of Georg Simmel, that 'the consciousness of constituting with others a unity is, actually all there is to this unity' and that, as such, 'Human beings X, Y, and Z constitute a collectivity (social group) if and only if each correctly thinks of himself and the others, taken together, as 'us' or 'we'' (1989, 146–147). So, for Gilbert, in order to form a plural subject, the members of the group need to view the group as a unified whole and act on that basis. But what does it mean for them to recognise such a unity and act it out? What is the principle of unification in question in terms of which they see themselves as unified? If the principle she has in mind is one of subjecthood, that they see themselves as a unified subject, then the discussion of subjective self-awareness calls into question the idea that such subjecthood could be enacted. If this is not the principle she has in mind, then she needs to spell out a (non-circular) alternative. Otherwise the idea that there can be unified plural subjects who are not subjects in the same sense as individuals is flawed at a quite basic level.

IV: Schmid

So far three attempts to defend the possibility of plural subjects without recourse to plural subjective self-awareness have been considered, and it has been found that are all flawed in one way or another. Let us turn, therefore, to the prospect of defending just such an account, that is, one which suggests that there is, or can be, such a thing as plural subjective self-awareness. As suggested above, this is a possibility that would be immediately discounted by most as a viable option on the grounds that it involves ascribing to groups precisely those properties that it does not appear to be even faintly possible that they might possess. Specifically, it would appear to involve the ascription of phenomenal self-awareness to groups. After all, it was suggested above that subjective self-awareness is a feature of the mode of presentation of experience; that we are subjectively self-aware in virtue of the fact that experience is necessarily presented as first-personal. Consequently, to say that groups can be subjectively self-aware would involve suggesting that they can have experience that is presented to them, as a group, as first-personal. But the suggestion that groups can have such experiences appears to affirm the kind of position that contemporary accounts of collective thought and action have actively eschewed. Rather, the thought goes, it must be recognised that phenomenal self-awareness requires, at the very least, something like a brain, and only individuals have brains, hence only individuals can be phenomenally self-aware.

This objection, however, commits an ontological fallacy in light of the account of subjective self-awareness defended thus far. It assumes that because phenomenal self-awareness requires a brain, that the subject of such awareness is settled by the brain in which it occurs. This is an assumption that is flatly denied by the account above, which argues instead that the subject of self-awareness is settled by the mode of presentation of that awareness. This is to say, that there is an individual subject just in case there is experience that presents itself as for that particular subject. Moreover, it is also an assumption that is shown to be false by the possibility presented by DID, in that it is perfectly reasonable to interpret an individual human suffering from such a disorder as embodying several distinct persons or subjects. In which case we have one brain with numerous distinct, and phenomenally self-aware, subjects. Of course, this particular interpretation might be brought into question, but the question of whether or not this is an accurate description of the disorder is distinct from the question of whether such a situation is conceivable, which it is.

Consequently, it would seem that a defence of plural subjective self-awareness does not have to make quite as radical a claim as it first appeared. On the one hand it does indeed have to claim that groups can have phenomenal experiences as groups and, therefore, be self-aware of themselves as a plural subject. On the other hand, all that needs to be shown to demonstrate that this is possible,

is that there is experience with a plural subjective character. This is to say, what stands in need of demonstration is that there is or can be experience that contains an intrinsic *for-us-ness* or whose mode of presentation is not first-personal, but *first-personal-plural*.

Precisely this claim has recently been defended by Hans Bernhard Schmid (2014b). He too accepts what looks like a broadly Zahavian account of subjective self-awareness. Namely, that there is a distinction between reflective forms of self-awareness, in which the subject is thematised and understood in terms of a concrete set of properties, and pre-reflective self-awareness, which is unthematic and non-ascriptive. And it is this latter form of self-awareness that he believes can be held in a plural form. In other words, he argues that there is a great deal of experience that is presented not as *for-me*, but as *for-us*.

This being said, defending the idea that some experience is presented in a uniquely plural form does present a challenge. In the case of individual subjects, it would be fair to say that a great deal of the argument in favour of subjective self-awareness rested on it being a hard, or perhaps impossible, feature of experience to deny. To deny that there is individual subjective self-awareness would be to deny that experience is presented as first-personal and this just seems extremely counter-intuitive. However, it is clearly not as counter-intuitive to deny that there is experience that presents itself as first-personal-plural.

Schmid's solution to this difficulty is to appeal to a combination of intuition regarding the nature of our social experiences, as well as what you might call suggestive evidence. His argument begins by identifying three features of a formally unified mind that he believes would not be possible without subjective self-awareness, respectively labelled *ownership*, *perspective*, and *commitment*. The first, ownership, refers to the way in which experiences or mental states are always presented as someone's. Thus a mind, understood as a totality of such experiences or states, is always someone's mind (2014b, 16). In other words, ownership refers to the intrinsic *for-me-ness* of experience. The second feature of a formally unified mind, perspective, refers to the way in which subjective self-awareness plays a role in distinguishing oneself from the world. It is only via subjective self-awareness, as an awareness of our own mind as ours, that we are able to conceive of that experience as perspectival in the sense that it involves an intrinsic distinction between perceiver and perceived (2014b, 16–17). Finally, the idea that subjective self-awareness explains the commitment exhibited by formally unified minds, affirms the point that Rovane attempted to reject. Namely, that experiential awareness of self entails a commitment to rational unity. In other words, for Schmid, subjective self-awareness is the reason why we are committed in any way to the content of our mental states:

‘Why should anybody bother whether or not the attitudes that happen to be in his or her mind are consistent? ... The reason is that the attitudes a person happens to find in mind are presented to her *as hers* ... To be aware of one’s belief that *p* is to see oneself *committed to the truth of p* and is therefore incompatible with the judgment that non-*p*.’ (2014b, 17)

For Schmid, therefore, unified minds must be understood in terms of these three features, and these three features can only be accounted for in terms of subjective self-awareness:

‘There may be a sense in which the mind is integrated on a purely ontological level, in terms of all of one’s attitudes being realized in one’s brain, or some such. But this purely ontological integration does not account for selfhood, a first-person perspective, and proper agency. The form of mental integration required for a self being related to the world in an active way requires self-awareness. Self-awareness is being aware of one’s attitudes *as one’s own*, as attitudes that are one’s own perspective *on* something, and as one’s own *commitments*.’ (2014b, 17–18)

Bearing this in mind, if it can be shown that groups also exhibit these three features, this would seem sufficient to demonstrate that groups can count as having unified minds as well as the subjective self-awareness that makes such unified minds possible. In light of this, Schmid suggests that there is evidence in favour of the idea that groups do in fact exhibit these features, or at least that such an attribution is not too farfetched, in the form of *common ownership*, *shared perspective*, and *joint commitment*.

Common ownership is essentially the idea, already discussed, that experiences are capable of not just being presented as for-me, but are also often presented as for-us. Interestingly, Schmid’s argument in favour of this possibility appeals to the idea of collective intention:

‘For example, it may not just be the case that I think that your intention to go for a walk is a good idea, and that I form the same private intention, but that we, together, come to intend to go for a walk. Similarly, a member of a congregation may have a belief that *p* not just in terms of some private cognitive attitude that he or she has taken over from others, but as something that involves his or her self-understanding as a member of the congregation.’ (2014b, 18)

While Schmid has, like myself, been highly critical of the intentionalist approach in general (accusing both content and mode based accounts of a form of vicious circularity), he clearly accepts the idea that groups are capable of collectively intending and that this capacity suggests that they must also be capable of having intentional episodes with a plural subjective character.

Once this is accepted, Schmid argues that the common ownership of experience naturally involves the existence of a shared perspective. As he puts it, it presupposes that, ‘There is a difference between how “we”, together, look at things, and the things as they are.’ (2014, 18) Which in turn involves a distinction between the experience and the object of that experience, or between

perceiver and perceived. To think of a certain mental state or experience as ours, is to think of *us* as a subject with a certain perspective on the world.

Finally, Schmid suggests that there is *prima facie* evidence in favour of the idea that we are naturally predisposed towards joint commitment:

‘Imagine the following case. At a party, you approach a person and try to strike up a conversation by saying something nice about the decoration of the room. Your conversation partner tells you that she hates it. Even though the topic of conversation does not really matter, as you don’t care about the quality of the decoration of the room, this will certainly bother you. Why do people have so much difficulty accepting that their own beliefs and value judgements are different from, and sometimes incompatible with other people’s, even under circumstances in which this is utterly inconsequential, such as in small talk? Why is it so important for people to agree with each other?’ (2014b, 19)

Schmid’s suggestion is that the reason for this need to agree is because there is a *normative pressure* towards collective coherence. We are naturally bothered by inconsistency between one another’s beliefs in the same way that we are pushed to resolve inconsistencies amongst our own beliefs. In the individual case this is because we view them as our own, and in the collective or social case this is because ‘we tend to be aware of our views as shareable and as individual contributions to a common perspective’ (2014b, 19).

Overall, therefore, Schmid believes that there is compelling evidence that groups exhibit the three features of a formally unified mind. And, if this is the case, then it is also the case that groups must have a form of subjective self-awareness. Schmid might not have quite done enough to definitively prove that this is the case, but he certainly presents evidence that makes it seem more reasonable.

Now, it is worth emphasising at this point that in terms of the broad strokes of his position, I am in complete agreement with Schmid’s account. I too believe that the only way to successfully defend the existence of plural subjects is to argue for the possibility of plural subjective self-awareness. This much is shown by the failures of the previous three positions considered. In turn such plural self-awareness requires a shift in the subjective character of experience and I too am convinced by the thought that sometimes experiences are presented as *for-us* rather than as *for-me*.

However, there are numerous problems with the way in which Schmid states his defence of plural subjective self-awareness. To begin with it seems that many aspects of his argument are simply confused. For instance, it is wrong to think of the ownership and perspective of mind as distinct phenomena once self-awareness is understood in the Zahavian sense of *for-me-ness*. Rather, they ought to be understood as two aspects of the same thing. To say that a mental state is owned is to say that it is presented as *for-me*, and to say that it is presented as *for-me* is to say that it forms part

of a perspective on the world; on this view of subjective self-awareness, ownership and perspective go hand-in-hand. As such, it is just confusing to distinguish between the two in the way that Schmid does. Furthermore, Schmid describes the relation between subjective self-awareness and the three features of a unified mind as a constitutive one: that subjective self-awareness ‘plays a constitutive role for the sort of unity in virtue of which the mind is somebody’s mind.’ (2014b, 16) But in the case of the ownership and perspective of mental states, such a relation cannot hold in virtue of the fact that it posits a distinction that does not exist. Subjective self-awareness should not be thought of as a distinct phenomenon from the intrinsic ownership and perspective of experience. Subjective self-awareness is simply synonymous with the fact that experience is presented with an intrinsic for-me-ness and, as such, as perspectival. It is not some extra entity over and above those features of experience – it just is those features of experience.

Consequently, in order to demonstrate that there is plural subjective self-awareness it would be sufficient simply to show that there are experiences that are owned in common or experienced from a shared perspective; either would amount to the claim that such experiences are presented as for-us and, therefore, have a plural subjective character. And yet Schmid’s actual argument in favour of common ownership and shared perspective does not go much further than an affirmation of the claim that attitudes are often presented as for-us, which, as suggested, really only amounts to saying the same thing in a different way. His only attempt to provide evidence in favour of this assertion is by appealing to the idea of collectively intending. But it is the very idea that groups are capable of collectively intending that many are still sceptical about. Moreover, part of the purpose of defending an account of plural subjects is presumably to be able to show that acts like collectively intending are possible, hence to appeal to them as part of one’s explanation is inherently circular.

The only real evidence that Schmid seems to discuss in favour of plural subjective self-awareness, therefore, is the idea of being jointly committed. If groups do exhibit the traits of being jointly committed to harmony amongst their various attitudes, then this would at least be suggestive of their having a form of plural subjective self-awareness as the ground of that commitment. However, the evidence that Schmid appeals to is wholly open to alternative interpretations. That we find it particularly disagreeable to disagree is a phenomenon that has not exactly gone unnoticed and is open to a host of, more individualistic, explanations. It is quite a leap to suggest that being annoyed by another person not holding the same opinion as you must be evidence in favour of a normative pressure towards coherence. For instance, what would this say about those contrarian fellows (not rare within the philosophical community) who enjoy disagreeing for disagreeing’s sake? And, perhaps much more importantly, what does it say about the possibility of building a society along

pluralistic lines? It seems, therefore, that Schmid is relying on fairly thin evidence to support his claim that individuals naturally feel a normative pressure towards coherence. And, even if he were successfully in establishing this claim, it would still only be suggestive of the more radical idea that such normative pressure is a consequence of thinking of our views as 'contributions to a common perspective'.

However, in many ways these brief critical points are much of an aside. Far more importantly for present purposes, Schmid's position brings to the surface a much more serious and immediate worry; one that presents a problem not just for the possibility of plural subjective self-awareness, but which brings in to question the very idea that any form of subjective self-awareness could be constitutive of subjecthood. The problem relates to how the subject given by subjective self-awareness is individuated. This is to say, it raises the question of how such a subject can be thought of as a particular individual at all, rather than some sort of universal consciousness.

This is quite a vague and confusing problem, so let us flesh it out a bit. Somewhat strangely, this difficulty comes to light in virtue of another error that Schmid makes regarding the nature of subjective self-awareness, in that he often describes it as a self-referential attitude:

'Self-awareness is... a self-referential attitude that each individual has to him- or herself and only to him- or herself.' (2014b, 17)

'In contrast to self-reflection, self-awareness is self-referential in an "immediate" and non-inferential way: it is not an awareness of self that, by virtue of being an object of consideration rather than a subject, is different from the subject of the act. Self-awareness is criterially empty because in self-awareness there is no question of whether the self to which I'm referring is me or not' (2014b, 13).

To begin with, it might strike one as odd to refer to a form of self-awareness as *self-referential*. Usually we understand referring to something as doing so by means of a word or sentence, which is obviously not the kind of reference that Schmid has in mind. However, let us assume that by 'reference' Schmid simply means something like 'picks out' or 'points to', which can be achieved without the use of language. Allowing this, it might still strike one as strange to think of self-awareness as *picking out*, or *pointing to*, the subject. For, this seems to imply that there is a pre-existent self or subject to which such self-awareness can point, and this is exactly what is denied on the Zahavian account of subjective self-awareness. The subject or self just is the self-givenness of experience and in its most minimal form it does not go beyond this givenness. What it is not is something in addition to that aspect of experience, for such a view would be open to the criticisms

considered in detail in the previous chapter⁴⁹. However, once again there is a charitable interpretation available, that views the kind of referring that Schmid is talking about as that which can bring into being the thing indicated; it is the very act of indication itself that allows us to understand the for-me-ness of experience as being constitutive of a subject.

However, even if one accepts all of this (and there may well be good reasons not to), the more fundamental problem with thinking of subjective self-awareness in this self-referential way, is that it allows Schmid to gloss over a very specific yet important problem. Namely, when he claims that an experience can be characterised by a for-us-ness, how is it the case that this plural subjective character designates the correct group of people? Presumably it must pick out a specific set of people, because surely the particular individuals that are members of a group constitute an essential or defining property of that group. Otherwise what would make it the group in question, rather than some other group? By taking subjective self-awareness to be a self-referential attitude, Schmid does not need to worry about this problem. He can say that the awareness in question picks out the correct set of people because the 'us' of for-us-ness *refers* to that specific set. If Jane, Jack, and myself constitute the plural subject of an experiential state, this is to say, on Schmid's account, that our experience is characterised by an intrinsic for-us-ness, where that 'us' refers to Jane, Jack, and myself. However, how is it possible that plural subjective self-awareness might refer in the manner? Crucially, it has been argued above that the subjecthood provided by subjective self-awareness must be understood as non-ascriptive in character. Bearing this in mind, how are the particular members of a group supposed to be included within the subjective mode of presentation of the experience without that experience becoming ascriptive and, therefore, an object taking form of awareness? By describing subjective self-awareness as a self-referential attitude, Schmid is able to mask an implicit slide that he makes between the two forms of awareness; by attributing to subjective self-awareness some of the features that must be reserved for objective forms of awareness, he can smuggle in the idea that plural subjective self-awareness might pick out the relevant members of the group.

However, having realised this problem for the plural case, it must be recognised that a similar problem also arises in the individual case. For, it has just been pointed out that subjective self-awareness involves the ascription of no properties whatsoever. It is, as such, a raw and abstract sense of self that arises only in virtue of the mode of presentation of experience. It is also, in virtue

⁴⁹ It is perhaps quite telling that Schmid at one point even aligns the view of self-awareness in question with the Kantian 'I think' that can accompany all of one's representations. In doing so he demonstrates that he is not aware of the subtle yet crucial differences between the Kantian and Zahavian accounts that are elucidated above.

of this fact, non-referential. How, then, can it also be first-personal? This is to say, if subjective self-awareness is neither a referential attitude nor does it involve the ascription of any properties, how is it able to pick out a particular individual at all? Why is the subject, that such self-awareness gives rise to, a particular individual subject?

This will undoubtedly strike most as a very strange question to ask. It seems strange because when we talk about experience it is so natural to think of that experience as already possessed by an individual. However, Zahavi's account of subjecthood already challenges this intuitive view of things. To be a subject is to be subjectively self-aware, where subjective self-awareness is an intrinsic feature of experience. Thus, there is no ontological subjecthood prior to experience and the matter of *to whom* an experience belongs is only settled by features intrinsic to that experience. But the subjective aspect of an experience is abstract and non-ascriptive. It, therefore, contains no features that would designate a particular individual, precisely because there is no such individual to pick out prior to the experience.

So, and this is the heart of the matter, it must be asked what it is that makes my experience *mine* and your experience *yours*? Qualitatively the subjective aspect of our experience is identical. Further, an appeal cannot be made to any feature external to experience in order to answer this question, even though this is how we would naturally go about such a task, for any such appeal will, of course, be circular. As has already been discussed, it cannot be claimed that it is because *my* experiences occur in *my* brain and *your* experiences occur in *your* brain because such an appeal already assumes that there is a brain that is mine and a brain that is yours. It also cannot be claimed, as Schmid seems to want to at times, that my experiences are those that are only accessible to me, because once again that presumes that there is already a self to do the accessing. The only way that this question could be answered is by appealing to some feature internal to the subjective aspect of experience, but nothing that has been claimed with regards to this aspect, so far, would suffice to achieve this role. The discussion of Zahavi's account of subjecthood in Chapter 3 made the quite intuitive assumption that subjective self-awareness is naturally individuating, but this assumption has been found wanting⁵⁰.

⁵⁰ It does seem that, in his most recent work, Zahavi recognises and dismisses something like this problem (2015, 22–25). However, I am not entirely convinced by this response and I take the argument presented in Chapter 5 to provide a more substantive solution; one which Zahavi himself appears to accept at various points, although not directly as a response to the problem of individuation.

V: Conclusions

In total, four attempts to defend the possibility of plural subjects have been considered and, in one way or another, each has presented difficulties. With respect to Bacharach and Rovane's accounts, this was because they both based their positions on a view of individual subjecthood that did not properly take into account the role of subjective self-awareness. For Gilbert's position, on the other hand, it was not because she denied this role of subjective self-awareness, but because her account of plural subjects attempted to avoid recourse to any subjective concepts in this sense. The result of which was that her position has difficulty with expressing the principle of unification on which the plural subject is based. Finally, Schmid's account faced problems, not because he denied that individuals are subjectively self-aware, or that such self-awareness could be possessed by groups, but because his own view of subjective self-awareness is flawed in such a way that it allows him to, unjustifiably, avoid a major obstacle that stands in the way of the possibility of plural forms of subjective self-awareness.

The insufficiency of Schmid's account, however, merely brings to light the fact that the concept of subjective self-awareness employed so far is itself flawed. Not only does it fail, as it stands, to ground the possibility of plural subjects, but it is not even clear that it can ground the possibility of individual subjects. The remaining chapters of this thesis, therefore, face a dual task. Firstly, the concept of subjective self-awareness must be returned to, in order to supplement it in such a way that it can be defended as an individuating form of self-awareness, without which it is hard to understand the concept's utility with regards to an account of subjecthood. Secondly, these repairs need to keep one eye firmly on the plural, for the primary concern of this thesis is a defence of ontologically collective forms of thought and action. Subjective self-awareness must be individuating, but it must not be so individuating that the individuals in question can only be understood as hermetically sealed units. In other words, what is sought is an account of subjecthood that leaves open the possibility that this could be achieved at non-individual levels.

Chapter 5

Bodily Self-Awareness

The attempt to defend and explain the possibility of ontologically collective forms of thought and action seems, at this point, to have reached an impasse. In Chapter 2 it was asserted that such an explanation would, owing to the intrinsic faults discovered in alternative approaches, require an account of subjecthood. In Chapter 3 just such an account was developed, but when this account was applied to the plural case in Chapter 4 a dead-end was reached. Specifically, if the intention is to retain the idea that subjecthood is constituted by subjective self-awareness, an account of the latter needs to be provided that explains how it might successfully individuate a particular subject, and it needs to do so while retaining the Zahavian idea that such awareness is pre-reflective and non-ascriptive. Moreover, as the ultimate aim of this thesis is to provide an explanation of thinking and acting together that appeals to the existence of plural subjects, it needs to be shown how such an individuating subjective self-awareness might also be possessed by groups.

For the time being let us put the second of these two issues to one side and focus first on the individuation of the subject. The broad aim of this chapter is to argue that subjective self-awareness can be thought of as individuating by taking into account the role of the body. So far the discussion of the subject has been had at a fairly abstract level that may capture important aspects of the self understood as an ontologically distinct entity, but which does not capture the way in which the self must also be understood as such an entity related to a world of objects. The significance of the body is that, on most philosophical accounts, it is the basis of this relation; the means by which a subject finds itself *in the world*. And the suggestion below will be that it is this significance that is the basis for how subjective self-awareness can be captured as an individuating form of awareness.

However, while there is a certain level of agreement regarding the fact that the body plays some role in the relation between a subject and the world, there is very little agreement regarding the nature of this role. In fact, the general discussion of the relation between self and body is guided by a rather frustrating multiplier effect, whereby the complexities involved in accounting for the role of the body are intensified by the countless ways there are to give an account of subjecthood in the first place. Consequently, the following discussion should not be viewed as anything approaching a comprehensive overview of the literature. Rather, it is a discussion guided completely by the

account of subjecthood that has been advanced thus far and the assumptions that this account forces us to make.

Bearing this account in mind, and specifically the distinction between subjective and objective self-awareness (where the latter is understood as reflective and the former pre-reflective), three broad ways in which the relation between the subject and the body might be accounted for will be considered. The first (section I), which will be referred to as the *subject-as-embodied* view⁵¹, is to argue that the self is related to a body through objective self-awareness. This is to say, that when I am aware of myself as a particular entity to which certain properties can be attributed (i.e. when I take a reflective stance towards myself), at least some of these properties will relate to my physical body. In other words, while there is a distinction to be made between the self and the body, in that they are not to be thought of as synonymous in any sense, they are nonetheless intrinsically linked in such a way that bodily-awareness can be thought of as a form of self-awareness. To those, like myself, naturally inclined against dualism, this may seem to be somewhat of a truism, but it should be noted that it does have to overcome a fairly longstanding Cartesian intuition:

‘Descartes thought that he was distinct from his body and could exist without it. The self that is distinct from its body is, according to Descartes, an immaterial substance. This immaterial self possesses a body and is so intimately conjoined with its body that it forms a union with it. The relation between self and the body is, Descartes insists, unlike the relation between a pilot and a vessel. If one were in one’s body like a pilot in a vessel one would not feel pain when one’s body is hurt. Nevertheless, the fact remains that each of us is, strictly speaking, distinct from his or her body. The self is a thinking and unextended thing. The body is an extended and unthinking thing.’ (Cassam 2013, 139)

For Descartes, the body is an object that is uniquely linked to the self, but it is an object nonetheless and, therefore, distinct from the immaterial self. Contrary to this, to suggest that bodily awareness can be a form of objective self-awareness is to say that, even though Descartes was correct to say that the body is an object, this does not exclude it from the category of selfhood. Rather, to be aware of one’s body is to be aware that one is a particular kind of physical being.

In order to defend this thought we will turn to an argument made by Gareth Evans, which has subsequently been built upon by many contemporary philosophers in the analytic tradition. This argument, broadly understood, is that certain forms of bodily awareness can count as self-awareness on the grounds that such forms of awareness enjoy a similar status to the seemingly privileged access we enjoy to our mental states. In other words, if awareness of one’s mental states counts as a form

⁵¹ It should be noted that the terminology used to distinguish the various accounts here, is often used in subtly different ways by other authors. The usage is closest to that adopted by Dorothee Legrand (e.g. 2013), but is not identical. For instance, she would include a Cartesian position as an embodied view of the subject, whereas I do not.

of self-awareness, then there is no reason to exclude awareness of one's physical states. However, while this line of thought is undoubtedly significant, and in many respects accurate, it will be suggested that it is not sufficient to provide a solution to the problem at hand, namely, the individuation of the subject. Specifically, it will be argued that if subjective self-awareness is not individuating in its own right, then there is nothing that can be said regarding objective self-awareness that would achieve such an individuation.

Thus, we will turn to a more radical option (section II), ubiquitous in the phenomenological tradition, which suggests that subjective self-awareness must itself already be understood as intrinsically tied to the body. Where the first view claims that the subject is embodied, in the sense that the physical body forms part of its broader self-conception, the second claims that the subject in even the most minimal sense is itself already *bodily* in its own right. This is to say, that if awareness of oneself as subject is presented in the pre-reflective subjective character of our experiences (for-me-ness), this subjective character must be understood as already bodily in some sense. If the subject-as-embodied view has to overcome a Cartesian intuition, this *subject-as-bodily* view has to overcome an even greater one. Where the former only argues that the body, as a physical object, can be understood as a part of the self, the latter argues that the presumption of the physicality of the body, and conversely the non-physicality of the subject, is itself a mistake.

However, while it will be argued that this view does successfully individuate the subject in a way that the subject-as-embodied view does not, it will also be suggested that this position leaves us with a problem. Specifically, it either leaves us with too thin an account of the bodily subject, one which does not do justice to many important phenomenological insights on the topic, or it transgresses the very principle that led to the current problem in the first place, that is, the non-ascriptive nature of subjective self-awareness. In short, it forces us to confront an issue that has been ignored thus far, in the form of the relation between subjective and objective self-awareness and how these two forms of awareness can be thought of as relating to one and the same self.

As a result, we will then turn to the third and final option (section III) in the form of a position motivated by the work of Maurice Merleau-Ponty's theory of perceptual experience. Merleau-Ponty's theory of perception not only asserts the intrinsically bodily character of the subject, but it allows us to recognise that any sharp division between subjective and objective forms of self-awareness rests itself on a faulty conception of experience. Rather, Merleau-Ponty argues that all intentional experience must be understood as an holistic structure involving figure (intentional object), background (the world), and the bodily subject. In many ways, this provides us with a way in which to harmonise the first two accounts of the relation between the subject and the body, as it

entails that we ought to view bodily self-awareness as a kind of *spectrum*. It is the very fact that all self-awareness is necessarily bodily, and that there is no sharp division to be made between intentional forms of awareness and pre-reflective self-awareness, that allows us to view both subjective and objective forms of self-awareness as part of a single continuum. It is only this account, in which our experiences of, and with, our bodies must always be understood as somewhere on a range from the more subjective to the more objective that can both make sense of subjective self-awareness as individuating, while retaining the idea that at some level such self-awareness is non-ascriptive, and account for the relation between subjective and objective forms of self-experience.

I: The Subject-as-Embodied View

A convincing argument to this end will require, however, a serious consideration of the alternatives. So, let us begin with the original challenge to Cartesianism and the idea that the subject is individuated through embodiment, in the sense that bodily awareness can count as a form of objective self-awareness. Taken in this sense, one's minimal subjecthood is still to be thought of as abstract and non-individuating in its own right, but it is nonetheless individuated (particularised in the world) in virtue of the fact that it is necessarily an embodied subjecthood⁵². Indeed, as will be seen, there are good reasons to accept that the subject is embodied in something like this sense. However, as will also be shown, there are also good reasons to be sceptical of the idea that such embodiment could be the basis for the individuation of the subject with which we are presently concerned.

To see why, let us consider how this claim might be defended. While there is no-one that defends such an account in precisely these terms (as this discussion presupposes a broadly Zahavian conception of subjecthood), over the past few decades a number of, predominantly analytic, philosophers have sought to demonstrate that the subject must be understood as embodied in some sense (correcting what was until that point, a reasonably large lacuna in the analytic tradition). The predominant thought in this regard has been to argue that awareness of one's body can count as a form of self-awareness in light of the fact that our awareness of certain physical properties enjoys the same status as our awareness of properties traditionally construed as mental. Thus, if our mental life is held to form a part of our selfhood, so the thought goes, our physical life must do so too. As

⁵² It seems that at least on some versions of the embodied account, it will remain a conceptual possibility that there might be such a thing as a disembodied subject (i.e. a sort of pure subjective self-awareness). However, this possibility ought to be of little concern. For, such a subject would be fundamentally cut off from the world that we must be concerned with, as it would lack its entrance onto that world. In other words, the question of whether there could be disembodied subjects is not a question for human beings, who must bear some relation to a body, even if that relation is a mere potentiality, in order to exist in a realm of other humans.

Evans puts it, 'If there is to be a division between the mental and the physical, it is a division which is spanned by the Ideas we have of ourselves.' (2001, 102)

To this end the most common approach, and Evans's own, is to appeal to the idea of immunity to error through misidentification. As discussed in Chapter 3, a judgement exhibits such immunity just in case it could not be the case that the maker of the judgement is mistaken as to whom the judgement applies. Thus, it cannot be the case that I come to be aware that someone believes that it is raining outside, take that person to be myself, and be wrong. While it was argued above that such a feature of self-reference is to explained in terms of the pre-reflective, and non-ascriptive, subjective character of experience, the pertinent thought with regard to present purposes is that such immunity is indicative of a judgement being founded on or provided by an information source that can only provide information about one's self and is only open to oneself. In other words, the idea is that such immunity is the primary indicator that the object of a particular intentional experience is oneself; if I have some means of gaining information that does not leave it open to any question as to whom that information relates, this would appear to indicate that the means in question bear a unique relation to oneself (i.e. that such a way of gaining information is a form of self-awareness). For, after all, I have no way of gaining information about anyone else in such a way that does not leave the identity of the person open to question⁵³.

For the sake of simplicity, let us call these special ways we have of gaining information about ourselves, forms of privileged access. The thought being that privileged access provides us with objective self-awareness, that is, awareness of oneself in terms of the particular properties that constitute one as a concrete and particular subject. The question then is whether or not we enjoy such privileged access to physical states as well as those traditionally classed as mental? Shoemaker, in elucidating his thoughts regarding immunity to error through misidentification, argues that we do not (1968). For Shoemaker, the only true form of privileged access is found in introspection, where introspection allows us unique access to our mental states and only *our* mental states. Gareth Evans, on the other hand, rejects this claim:

'Unfortunately, many philosophers give the quite mistaken impression that it is only our satisfaction of mental properties which gives rise to judgements exhibiting immunity to error through misidentification. This is tantamount to the claim that self-conscious thought rests only upon knowledge we have of ourselves as mental or spiritual beings. And this in turn generates the unfortunate, and quite inaccurate, impression that in thinking of oneself self-consciously, one is paradigmatically thinking about oneself as the bearer of mental properties,

⁵³ For a more detailed and nuanced defence of the link between immunity to error and self-awareness see Bermúdez (2013).

or as a mind—so that our ‘I’ thoughts leave it open, as a possibility, that we are perhaps *nothing* but a mind.’ (2001, 105–106)⁵⁴

In light of this, Evans outlines two broad categories of bodily awareness that give rise to immune self-judgements. The first he describes as ‘a general capacity to perceive our own bodies’ (2001, 109). It is this general capacity that has undoubtedly been the dominant focus of subsequent accounts of embodiment, but the actual mechanisms that are usually included within it vary enormously (partly in line with advances in scientific research in the area) and it is difficult to say in any precise terms what the key feature is that holds this grouping together. Nevertheless, a crude way of putting it is to say that it includes what one might call *internal bodily senses*, that is, all those ways we have of gaining information regarding our bodies from the inside (often referred to as *interoception*). The paradigmatic example of which, and by far the most discussed, is proprioception, where proprioception refers to a complex informational apparatus that provides us with awareness of the relative positioning of our limbs and the effort or force applied in movement (e.g. I do not need to look at my arm in order to know that it is swinging by my side as I walk or delivering a right hook in the course of a boxing match). Thus, as Evans points out:

‘None of the following utterances appears to make sense when the first component expresses knowledge gained in the appropriate way: ‘Someone’s legs are crossed, but is it my legs that are crossed?’; ‘Someone is hot and sticky, but is it I who am hot and sticky?’; ‘Someone is being pushed, but is it I who am being pushed?’ There just does not appear to be a gap between the subject’s having information (or appearing to have information), in the appropriate way, that the property of being *F* is instantiated, and his having information (or appearing to have information) that *he is F*.’ (2001, 109)

Evans’s point being that if I have a proprioceptive awareness that someone’s legs are crossed it is not necessary for me then to ask whose legs are in question. Part of what it is for the information to be provided by proprioception is for such questions to be irrelevant.

The second category, which has received far less attention, relates to the awareness we have of our spatio-temporal location in virtue of our perception of the world (often referred to as *exteroception*). A useful way of contrasting this with the first category, therefore, might be external, as opposed to internal, bodily awareness. The first is an awareness of the general condition of our bodies through internal sense, while the second is an awareness of things such as the location and orientation of our bodies through external sense:

‘Included here are such things as: knowing that one is in one’s bedroom by perceiving and recognizing the room and its contents; knowing that one is moving in a train by seeing the world slide by; knowing that there is a tree in front of one, or to the right or left, by seeing

⁵⁴ Although later Evans points out that once we become aware of the embodied nature of the subject, the very distinction between the mental and the physical begins to appear quite artificial (2001, 109).

it; and so on. Once again, none of the following utterances appears to make sense when the first component expresses knowledge gained in this way: 'Someone is in my bedroom, but is it I?'; 'Someone is moving, but is it I?'; 'Someone is standing in front of a tree, but is it I?' (Evans 2001, 111)

There appears to be an inference involved here that is not necessary with those forms of bodily awareness in the first category, but this inference makes it no less absurd to ask to whom the information in question relates; it is part of what it is to perceive the world, to be able to infer from this perception a judgement regarding one's body that is immune to error through misidentification.

The standard response to this line of thought is to argue that it seems as if our awareness of our physical states is capable of going awry in a way that our awareness of our mental states is not. For instance, take the proprioceptive awareness that my legs are crossed. Perhaps unbeknownst to me, a mad scientist has been tinkering with my body and managed to wire my brain up to someone else's body in such a way that I receive the proprioceptive signals of that body. In such a scenario it seems that I might think my legs are crossed when in fact it is this other person's legs. Surely this is a case in which there is an error due to my having identified the wrong subject?

Provided that we allow that the case is conceptually possible, there are several potential lines of response. One is to deny that an error occurs. In other words, one could argue that 'my legs' are just whatever legs I happen to be aware of in the proper way, such that if I have a proprioceptive awareness of a pair of legs not connected to the rest of my body, these are still my legs⁵⁵. However, this has the rather counter intuitive implication that my body, even understood as a determinate physical object, might not be continuous with a single biological organism. Another tactic might be to point out that one has to adopt a quite radical dualism to think that similar cases will not arise in the case of mental states. As long as one holds that one's mental states have some physical instantiation it appears conceptually possible that one's brain might be wired to someone else's in such a way that similar errors will occur (for instance, that whenever they form a certain belief I am immediately aware of this belief in just the same way that I would be if the belief was formed in my own brain).

However, Evans's own response is actually far simpler and, for that reason, more convincing. He suggests that there is clearly an error in such a case but the error has nothing to do with misidentification. Instead, my error concerns whether or not my legs are in fact crossed; due to receiving certain signals I believe that they are when they are not. It is as simple as that. It cannot

⁵⁵ Quassim Cassam describes this as the idealistic account of the body and attributes such a view to Locke (1998; 2013).

be that I misidentified the subject because it is absurd to believe that any identification was involved in the first place:

‘[W]e cannot think of the kinaesthetic and proprioceptive system as gaining *knowledge* of truths about the condition of a body which leaves the question of the identity of the body open. If the subject does not know that *he* has his legs bent (say) on this basis... then he does not know *anything* on this basis. (To judge that *someone* has his legs bent would be a wild shot in the dark)’ (Evans 2001, 110)

For a misidentification to have taken place the subject would have had to first come to an awareness, proprioceptively, that there was a pair of legs that was crossed. Upon realising this they would then have had to (mistakenly) identify the legs in question as their own (presumably on the basis that they *felt* a lot like their own, if such an idea can even be made sense of). That this is not a reasonable description of how a subject would come to form such a judgement should be sufficient to illustrate Evans’s point. Just because the signals happened to be coming from someone else’s body does not make the case differ in any way from that in which I just happen to be wrong about the state of my own legs (perhaps because they are numb from sitting on them for too long, or some such scenario). In other words, there are a whole host of reasons why I might form a mistaken judgement, on the basis of interoception, regarding my physical state. One of which is that a mad scientist has wired up my brain to receive the proprioceptive signals of someone else’s body. But none of these reasons demonstrate cases in which I have misidentified the *subject* to whom the judgement applies, because none of them involve a process of identification in the first place.

The nuances of this line of thought have been debated at length since Evans’s initial offering on the subject⁵⁶, but it is not necessary to discuss these nuances in any detail in order to see why this approach is not going to be sufficient with regards to the individuation of the subject. Most immediately, it should be realised that it involves an assumption that has actually already been rejected above (Chapter 4, section III). Specifically, it assumes that forms of privileged access, that is, informational channels that give rise to immune self-judgements, cannot fail to be about oneself. On the basis of this assumption, the embodied account accepts that privileged access entails self-awareness (in that any information gained in this privileged way would necessarily be about oneself). However, while this might be the case in all ordinary circumstances, a marginal case has already been considered where this claim is found to be untrue, in the form of Dissociative Identity Disorder (DID). Above an argument made by Carol Rovane was defended, in which she suggests that it is reasonable to understand an individual human being suffering from DID as embodying two or more

⁵⁶ Bermúdez (2013), for instance, offers a complex taxonomy of the various different ways that we can gain information regarding our bodies that goes well beyond Evans’s initial two categories. Conversely, for a particularly unique critique of the idea that bodily awareness is a form of self-awareness see Smith (2006).

distinct subjects (despite the fact that the basis on which she held this to be the case was found wanting). If true, or even if at least conceivable, this involves the suggestion that there can be two distinct subjects who have privileged access to one another's thoughts. For, as Rovane pointed out, it is well documented that alter personalities can be, and often are, co-conscious. As such, this appears to be a fairly reasonable case in which there might be privileged access that is not a form of self-awareness⁵⁷. Similarly, in the physical case we can point to pathological disorders such as Alien Hand Syndrome (AHD), whereby an individual will experience the actions of one of their limbs as outside of their control. What is peculiar about AHD is that while the person in question usually retains a complete awareness of the limb interoceptively, they experience the sensation gained in this way as not relating to themselves, often suggesting that the limb is being controlled by someone else. Once again this seems to be a clear case in which privileged access remains, but not as a form of self-awareness.

It might be objected that this is a mischaracterisation of the nature of experience in such cases. While the patient suffering from AHD experiences a hand that they do not take to be their own, it must still be the case that their interoceptive experience of the hand is presented as *their* experience. It is just a first-personal experience of someone else's hand. Similarly, in the case of DID, while alter personalities might be aware of one another's mental states, this awareness itself must be characterised as first personal. However, to think that this constitutes an objection is to miss the point slightly. Indeed, when alter A is aware, introspectively, of one of the beliefs of alter B, it is correct to assert that there must be something that is presented as first-personal for alter A. In this case, it will be the awareness of alter B's belief. Nevertheless, the significant point is the fact that what will *not* be presented as first-personal is alter B's belief itself. Conversely, in the case of alter A's own beliefs, it is this extra step, in the form of an *awareness of* the belief, that is missing. In other words, when it comes to alter A's own beliefs, it is these beliefs themselves, that is, the experience of believing, that is presented as first-personal. And it is this first-personal presentation of the beliefs themselves that marks them as those belonging to alter A.

What this puts paid to is the claim made by Bermúdez (2013, 167) that, 'If I know through introspection that someone is currently thinking about Southern Arizona, then I know that I am thinking about Southern Arizona.' It seems strange in the case of DID to think that there might be

⁵⁷ Interestingly, I am not aware of anyone who has considered whether such cases give rise to judgements that are immune to error through misidentification where the subject in question is not 'I'. For, it seems that *alter personality A* will be phenomenally conscious of the mental states of *alter personality B* in precisely the same way in which they are conscious of their own states. As such, there seems to be no obvious reason why judgements made on this basis shouldn't be considered immune to error through misidentification, as it is not the case that an identification has taken place: alter A just immediately knows that the mental state in question is that of alter B.

a person able to report directly on the experiences of another, but this is simply because it is so far removed from the way in which we ordinarily experience mental states and it is testament to how closely we associate experience with a first-personal presentation. But there is no reason to think that one alter cannot make a report of 'alter B believes that p' in precisely the same way that alter B can report 'I believe that p'. In other words, the guiding assumption at work in the subject-as-embodied view, that information gained through channels of privileged access necessarily relates to one's self, is found to be mistaken. Rather, if one accepts that what makes a particular experience one's own is its first-personal character, that is, the pre-reflective awareness of self that is part of all experience, one must reject the idea that there is any mode of access that is specific to the self. For, a mode of access is a way in which one can gain certain information, but no mode of access guarantees that the information gained will be presented as relating to oneself.

Nevertheless, it is important to note that the weight of this criticism is not aimed at the subject-as-embodied view offered by the likes of Evans, which can take this fact in its stride. The advocate of embodiment can readily admit that there may be particular circumstances, usually rare pathological cases, in which privileged access to physical states is not presented as first-personal, but point out that in all normal circumstances first-personal presentation will be part and parcel of privileged access. Thus, they can amend the position to claim that bodily awareness is a form of objective self-awareness just in case it is via some form of privileged access *and* that the information gained in this way is presented immediately as for-me. For, it should be noted that in the case of DID it was argued that precisely what made it possible to make sense of two subjects being co-conscious was the fact that those experiential states would only be presented as first-personal for one of the subjects in question.

Unfortunately, however, while this amounts to a small amendment in terms of the subject-as-embodied view, it is a significant one in terms of the problem of individuation. The problem currently at hand is the non-individuating character of subjective self-awareness; there is nothing contained within the subjective character of our experiences that entails the designation of a particular individual. The present suggestion is that the subject, constituted by subjective self-awareness at the most basic level, might be individuated in virtue of its relation to a particular physical body of which it is objectively self-aware. But on the understanding of subjecthood offered thus far, there can be no objective self-awareness independently of subjective self-awareness, regardless of whether the object in question is one's body or one's mind. What makes my awareness of the particular properties that can be ascribed to me, whether they be mental or physical, an awareness of myself, is the subjective character of that awareness, that is, whether or not they are

presented as for-me. And here we come back to the original problem: there is nothing that has been said about this pre-reflective sense of self, thus far, which would account for it as an individuating form of self-awareness.

This is not to suggest that all experience that is presented as for-me is a form of objective self-awareness. The vast majority of my intentional experiences take as their object features of the world that have nothing to do with my self. What the subject-as-embodied view recognises is that there are forms of intentional experience that are intimately and regularly tied to an awareness of oneself, and that these forms of privileged access relate not just to the mind but to the body as well. Nevertheless, that such forms of awareness, whether they relate to the mind or the body, cannot straightforwardly count as forms of *self*-awareness is demonstrated by the possibility of them going wrong in marginal cases. Thus, it must now be recognised that objective self-awareness has to be understood as a combination of privileged access *and* first-personal presentation (subjective self-awareness), and it is the addition of the latter that means that the subject-as-embodied view, for all that it gets right, is of no use when it comes to the individuation of the subject. As long as subjective self-awareness is in its own right non-individuating, it is irrelevant whether or not objective self-awareness is in part an awareness of a particular body as my own, for there is no particular individual subject to do the owning.

II: The Subject-as-Bodily View

Consequently, if the account of the subject that has been developed thus far is to hold up, it would seem that the individuation of the subject will have to be intrinsic to subjective self-awareness. Let us, therefore, turn to the suggestion that subjective self-awareness is already a form of bodily awareness. As Henry and Thompson put this claim:

‘[P]henomenologists from Husserl onward regard the original form of bodily experience not as consciousness of the body as intentional object but rather as an intransitive (non-object-directed) and prereflective bodily self-awareness.’ (2013, 239)

To put it differently, phenomenologists would argue that the above account of embodiment fails to properly describe all of the ways in which the body features in our experiences. Almost all of the examples in the literature on embodiment focus upon our *awareness of* our bodies, such as the much worn example of one’s knowledge that one’s legs are crossed (a proprioceptive awareness of the positioning of my legs wherein my legs are taken as intentional object). This is a consequence of the fairly intuitive presumption of the physicality of the body; the body is taken, first and foremost, as a physical object in the world and it is then asked, on this basis, whether this object can form part of the self. As such the question of the body in this context can only be the question of whether we

are aware of ourselves as a particular kind of object. However, it is precisely this assumption that is challenged by phenomenologists:

‘Obsessed with being, and forgetting the perspectivism of my experience, I henceforth treat my experience as an object and I deduce it from a relation among objects. I consider my body, which is my point of view upon the world, as one of the objects of that world. I repress the consciousness that I had of my gaze as a means of knowing and I treat my eyes as fragments of matter. From then on my eyes are placed within the same objective space where I attempt to situate the exterior object and I believe that the projection of the objects upon my retina brings about the perceived perspective.’ (Merleau-Ponty 2013, 99⁵⁸)

In other words, the presumption of the physicality of the body is shown to be mistaken once we distinguish subjective from objective self-awareness, which presents an alternative way in which to account for our experiences of the body. To this end, it is useful to return to Evans who actually goes some way to demonstrating how this can be the case. As outlined above, he distinguishes two broad ways in which one can have immune bodily self-awareness, where the first relates to those ways we have of knowing the general condition of our bodies through interoception, while the second is concerned with the awareness we have of our body’s spatio-temporal location on the basis of (an inference from) exteroception. Above these were presented as two parts of the same category: forms of self-awareness that take the body as object. However, while this point is not fully developed, Evans appears to have been quite aware that there is something significantly different about the second form of awareness:

‘The explanation of the importance of this way of gaining knowledge is not hard to find. Any thinker who has an idea of an objective spatial world—an idea of a world of objects and phenomena which can be perceived but which are not dependent on being perceived for their existence—must be able to think of his perception of the world as being simultaneously due to his position in the world, and to the condition of the world at that position. The very idea of a perceivable, objective, spatial world brings with it the idea of the subject as being in the world, with the course of his perceptions due to his changing position in the world and to the more or less stable way the world is. The idea that there is an objective world and the idea that the subject is somewhere cannot be separated, and where he is is given by what he can perceive.’ (Evans 2001, 112)⁵⁹

To put it simply, our perceptions of the world demand that we are in the world; to perceive *something* presupposes that we understand ourselves as *somewhere*. Now, while this argument is largely motivated by Evans’s attempts to defend a certain form of materialism, of which it is not

⁵⁸ All references to Merleau-Ponty’s *Phenomenology of Perception* use the original French pagination, as provided in the translation by Donald A. Landes.

⁵⁹ This is in large part a direct response to an argument made by Strawson, and it is perhaps even more illuminating that he states in a footnote shortly after: ‘Do we really have to go any further than this in order to answer Strawson’s questions “(1) Why are one’s states of consciousness ascribed to anything at all? And (2) Why are they ascribed to the very same thing as certain corporeal characteristics, a certain physical situation, etc.?”’ (Evans 2001, 112; the quote comes from Strawson (1964))

necessary to speak here, it is interesting how close this line of thought is to a claim that is ubiquitous in phenomenological thought. Dorothee Legrand describes this claim as the *constitutive embeddedness* of the subject (2013, 207–209) and suggests that it is both demanded by the *dual givenness of experience* and demands that we understand subjective self-awareness as intrinsically *bodily*.

Let us consider this double move in more detail. As has been discussed, the idea that subjective self-awareness is constitutive of subjecthood is based on a distinction between awareness of oneself as subject and awareness of oneself as object. In turn this thought is founded on the idea that all experience must be understood as encompassing two forms of awareness, and it is this necessary double form that Legrand refers to as its dual givenness. On the one hand all experiences are intentional in the sense that they are an awareness of something, that is, they take some feature of the world as their object. On the other hand, in taking an intentional object our experiences of the world must also be presented pre-reflectively as for someone. For, it is not possible to make sense of a free-floating experience not anchored to a subject. Allowing this, how then are we to make sense of this dual givenness? This is to say, what must be true of how experience presents itself to us such that it appears simultaneously as both *of something* and *for someone*? As Legrand points out, what must immediately be the case is that our experiences are presented as for a subject that is constitutively embedded in the world. Minimally speaking, we are nothing over and above the subjective character of our experiences, but these are experiences that are intentional and world-regarding. If, therefore, we are nothing over and above the subjective character of these experiences it must be the case that we are presented to ourselves as a being in the world. For, the experiences in which we are given to ourselves are always necessarily *of the world*. To put it differently, we are constitutively embedded in the sense that we are constituted, as minimal subjects, by world-regarding experience.

Having established this first move, one must then explain this appearance of the subject as intrinsically a being in the world. Subjective self-awareness is an embedded self-awareness, but what is the means of this embedding? And here it seems that the only available answer is to say that I am embedded in virtue of my bodiliness:

‘For the self to belong to the world, there is no other way than being corporeal. Not only being an experiencing subject, but more specifically being an experiencing *body* is necessary for there to be an experienced world at all.’ (Dorothee Legrand 2013, 209)

To assert that I am presented to myself as a being in the world is to assert that I am presented to myself as an intrinsically bodily subject. What other means of being in the world could possibly be open to me? This is not to say that I am presented as a subject of flesh and bone. To say this would

be to once again presuppose the physicality of the body, which is precisely what is not being assumed:

‘[M]y original body-awareness is not a type of object-consciousness, is not a perception of the body as an object, but a form of immediate, prereflective, self-awareness.’ (Zahavi 2002, 21)

It is not, therefore, that the body is first presented to me as a physical object which poses the problem of how to integrate this object with the self. Rather, the body is what we are at even the most minimal and pre-reflective level in the sense that experience presents us as located in the world. All forms of self-awareness are, therefore, necessarily forms of bodily self-awareness. To put it differently, on all accounts, even the Cartesian position which divorces the body from the self, the subject finds itself related to a world through a body. What the particular insight of phenomenology consists of is in illustrating how, once we understand the subject as constituted by world-regarding experience, both this embeddedness and bodiliness must be understood as already intrinsic features of one’s most fundamental or minimal subjecthood. The nature of experience, as a perspective on the world, presupposes that it is presented as for a subject located in that world, and this fact of being located must be understood as the most primordial form of bodiliness⁶⁰.

Returning once again to Evans’s account of bodily awareness via exteroception, it appears that one can now provide a more compelling explanation of the phenomenon. While Evans realises that there is something important about this form of bodily awareness, his mistake rests in assuming that this awareness must be of the body as an intentional object. Rather, to take the example of knowing that I am in my bedroom by perceiving the room, it appears that my primary experience is bodily in a much more fundamental way. It is an experience that takes the room as its object, but in which this very object taking would be incomprehensible if the experience were not understood as for an embedded subject whose means of being in the world is its intrinsic bodiliness. This is not to deny that we could take the reflective step in which the object of my experience ceases to be the room *per se* and instead becomes myself in the room, but myself taken as the object in this sense must be understood as distinct from the bodily subject for whom this determinate self appears. The point is not to refute the insights of Evans; we can take our bodies as object and in doing so they can form part of our objective self-awareness. The point is that being able to do so already rests on one’s

⁶⁰ It must be acknowledged that mere *spatial locatedness* is a particularly thin sense of *bodiliness*. It might, therefore, be argued that it is a bit misleading to refer to the subject constituted by such subjective self-awareness as *bodily*; being spatially located might be a necessary condition of being considered a bodily subject, but it is not sufficient. It is at least a conceptual possibility that a subject could be spatially located without being bodily in any substantial sense. However, as will be seen below, it is precisely this claim that is problematised by Merleau-Ponty’s account of the subject and bodiliness. As such, it makes sense to treat this locatedness as the most primordial form of being bodily for the time being, allowing that the pertinence of this terminology rests on an argument that is still to be made.

prior intrinsically bodily subjecthood. In this context it is only the fact that my experience of the room must come with a pre-reflective awareness of myself as the bodily subject of the experience that allows me, in a second move, to extrapolate from this experience to an awareness of my body as a physical object amongst other physical objects.

The intuitive objection to this claim is to point to potential cases in which there might be experience but in which this experience is divorced from any sense of locatedness. For instance, imagine a case of a subject placed in a sensory deprivation tank of some sort. Also imagine that this tank is placed in a zero gravity environment and that the subject in question has been given a powerful anaesthetic that removes any form of interoceptive awareness that they might have. And, just for good measure, let us imagine that the subject was knocked unconscious prior to being placed in the tank, such that they have no knowledge of where they are or how they got there. Would this be a case in which the subject has experience and is given to themselves as the pre-reflective subject of that experience, but in which any sense of bodiliness has been removed from the equation? Surely not. For a start, one might take the lengths that one had to go to deprive the subject of its pre-reflective sense of bodiliness as evidence of just how essential to everyday experience this bodiliness is. More importantly, however, it is the fact that the subject's reaction to such a scenario will be one of extreme disconcertion, which is truly telling. Indeed, the first thing that they are likely to do is to try and regain some sense of orientation, probably by attempting to move their body and seeing what happens. In other words, it will be precisely the fact that the subject will experience the scenario as a case of extreme *loss* that demonstrates the significance that its bodiliness has for it in the first place.

So, returning to the driving question, does this subject-as-bodily view achieve the desired individuation of the subject? It would certainly seem so. It is now clear that the nature of experience demands that the subject is understood in terms of a particular bodily perspective. The subject is thus individuated, or rather subjective self-awareness is constitutive of a particular individual subject, in virtue of the fact that one is given to oneself in experience as first and foremost a perspective on the world, distinct from the perspectives of others, and as bodily as the means of that perspective. As Zahavi aptly summarises with respect to Husserl's thoughts on the matter:

'Let us assume that I am sitting in a restaurant. I wish to begin to eat, and so I pick up the fork. But how can I do that? In order to pick up the fork, I need to know its position in relation to myself... On the dinner table, the perceived fork is to the left (of me), the perceived knife is to the right (of me), and the perceived plate and wineglass in front (of me). As Husserl writes, every perspectival appearance implies that the embodied perceiver is itself co-given as the zero point, the absolute indexical 'here' in relation to which every appearing object is oriented. As an experiencing, embodied subject I am the point of reference in

relation to which each and every of my perceptual objects are uniquely related. I am the centre around which and in relation to which egocentric space unfolds itself.' (Zahavi 2002, 19–20)

Thus, where the above exposition of the subject was concerned largely with the 'I', the essential role of the body informs us that an account of the subject must be equally concerned with the 'here'. It is not that experience is intrinsically presented with a characteristic for-me-ness, but with a *for-me-here-ness*⁶¹. The importance of this *here-ness* should not be understated. Where 'me' gives the impression of a separation of the subject from the world, 'here' indicates that the subject must be understood as in the world and, in virtue of this, particular. The point being that the nature of our experiences do not present themselves as for a subject divorced from the world, but for a subject in the world and, therefore, particularised in virtue of its relation to space. But neither the me-ness nor here-ness of experience should be taken as ascribing any substantive property to the subject. Rather, they must be understood as formal or structural features of all experience and they offer a view of subjective self-awareness that is still non-ascriptive in character. One might be able to infer from this character some determinate property about oneself, as Evans does in the case of his knowledge that he is in his room. But to do so would require the taking of a reflective step that goes beyond that which is given in pre-reflective self-awareness.

Nevertheless, as phenomenologists have also been concerned to make clear, experience does not present itself as for a subject that is *merely* in the world, and the sense in which the subject is located is the most minimal and thin sense in which the subject is bodily. While phenomenology in general draws a distinction between the subjective and the objective body, that is, between the study of the body as a lived subjecthood and the anatomist's study of flesh and bone, there is no phenomenologist who takes this distinction to reduce the former to a passive *here-ness* (as if everything else were the domain of a physical science). Rather, the most important insight of phenomenology's take on the bodily nature of the subject, is the idea that the subject must be understood as essentially active:

'Consciousness is originally not an "I think that," but rather an "I can." ... The gesture of reaching one's hand out toward an object contains a reference to the object, not as a representation, but as this highly determinate thing toward which we are thrown, next to which we are through anticipation, and which we haunt. Consciousness is being toward the thing through the intermediary of the body ... Motricity is thus not, as it were, a servant of consciousness, transporting the body to the point of space that we imagine beforehand. For us to be able to move our body toward an object, the object must first exist for it, and hence our body must not belong to the region of the "in-itself."' (Merleau-Ponty 2013, 171–173)

⁶¹ In fact, bearing in mind the temporal nature of the minimal subject (see Chapter 3 section IV), the most accurate description of the subjective character of experience would be as a *for-me-here-now-ness*. For, temporality must be understood to be just as fundamental to subjective self-awareness as spatiality.

In other words, the sense in which the subject is bodily is that in which it is not just in the world, but *toward the world*. For instance, to return to the example of sitting in a restaurant, it is not just that my perception of the various utensils unfolds in front of me in a way that presupposes my spatial relation to them, for I do not receive my perceptions as a static and passive image. The fork does not just appear to the left of me, but within reaching distance, as graspable by me, and with the various affordances that it has for being used by me (such as how it will feel in my hand or how effective it will be for eating with). And each of these features of my experience presuppose various features of my own pre-reflective self-understanding. The fork can only appear as within reaching distance for a subject of a certain shape and size and it can only appear as graspable for a subject with hands that are familiar to it. As such, the nature of experience demands that the bodily subject constituted by such experience is not to be understood as a passive receiver of information. Experience is shaped by our ability to move around the world and interact with it, and the bodily subject can only find itself through activity. As Zahavi puts it:

‘The body is not first given for us and subsequently used to investigate the world. The world is given to us as bodily investigated, and the body is revealed to us in its exploration of the world.’ (2002, 20)

As outlined above, Evans pointed out that we can infer from our experiences to an explicit thematic awareness of ourselves as bodily located. What he failed to recognise, however, was that the nature of our experiences allows us to infer from these experiences to many other features of our bodiliness; features that go beyond mere location. And, just as with spatio-temporal location, the fact that we can infer to these details shows that they are already there, in a sense, in the original experience, presupposed as part of the pre-reflective and unthematic subject of that experience.

In saying this, however, a contradiction or tension in the account appears to unveil itself. While the idea of a pre-reflective locatedness appeared sufficient to establish the individuation of the subject given as located, the current suggestion is that our pre-reflective awareness of self already goes well beyond this minimal and passive sense of being bodily. But there seems to be no way in which these insights might be incorporated into the account of the bodily subject without transgressing the problematic principle that motivated this turn to the body in the first place, namely, the non-ascriptive nature of subjective self-awareness. It was a core tenet of the above exposition of such self-awareness that one knows oneself first and foremost not as a particular kind of object, identified through its properties, but simply as the pre-reflective subject of experience, and it was the fact that subjective self-awareness must be understood in this way that led to the problem of the individuation of the subject. However, the account of the bodily subject offered by phenomenology is not one of a simple, abstract bodiliness or for-me-here-ness, which can be understood as non-

ascriptive. Instead, the point is that, even pre-reflectively, experience presents itself as for a bodily subject in a much richer and, most importantly, active sense. It presents itself as for a subject with properties such as size and shape, and with certain potentialities to act upon the world. The subjective character of our experiences, therefore, appears to go well beyond the simple and non-ascriptive. Recognising the essentially bodily character of the subject leaves us with a dilemma: how can one accept that the subject given in the pre-reflective subjective character of experience must be understood as an active, and more or less determinate, entity, while maintaining a view of minimal subjecthood as non-ascriptive? Is it possible to harmonise these features?

III: The Spectrum View

In order to see how such a harmonisation might be possible, it is necessary to leave the safe and relatively unambiguous terrain that has been mapped so far. For, I propose, the philosophical utility of the body is precisely the fact that its role in the subjective life of the individual cannot be captured unambiguously and it is this ambiguity that can heal the ruptures in the present account. To put it differently, we should neither reject the insights of the phenomenological take on the subjective body, nor should we discard the idea that subjective self-awareness must be understood as primarily non-ascriptive in some sense. Rather, it is the underlying view of experience, in which these are seen as mutually incompatible positions, that needs to be re-evaluated. As will be shown, Merleau-Ponty's theory of perceptual experience allows us to view pre-reflective forms of bodily self-awareness and object-taking intentional awareness, not as two distinct types of awareness, but as mutually reinforcing poles of the same experiential continuum. On this basis, subjective (pre-reflective and non-ascriptive) and objective (reflective and ascriptive) forms of bodily self-awareness cease to appear as wholly distinct senses of self, and instead appear as the two ends of a single spectrum. Indeed, it is the fact that the body is both subject and object, and yet perhaps never either completely, that allows us to properly account for the dual givenness of experience and unite objective and subjective self-awareness on the very same spectrum of being.

In order to grasp the import of Merleau-Ponty's account of experience, it is useful to begin with a consideration of an account that attempts to embrace this view of bodily self-awareness as a spectrum. As Dorothee Legrand has suggested, 'If the self is bodily in a non-metaphorical sense, then bodily self-consciousness always involves the intermingling of subjective and intentional aspects.' (2013, 213) To this end, Legrand asks us to consider the following four part experiential story (Ibid.):

(1) I see a rose within reaching distance.

- (2) I reach out, my hand moves towards the rose.
- (3) Getting closer, I remember that roses are delicate but thorns are sharp, I thus pay attention to the movement of my hand so that I damage neither the rose nor my hand. I touch the rose, squeeze its stem and feel a sudden pain at the tip of my finger.
- (4) I scrutinize the skin of my finger in the attempt to detect a hidden thorn.

For Legrand, these four stages correspond to four distinct ways in which we can possess bodily self-awareness, ranging from the most subjective to the most objective. Let us consider these four stages in more detail:

(1) When I see the rose within reaching distance I experience my body in the most subjective or pre-reflective sense. It is that sense which does not go much beyond the idea of for-me-here-ness, in that this simple perceptual experience requires very little in terms of my body understood as either a concrete entity or an active agent. With such perception we experience ourselves as oriented towards the rose and, therefore, in a way that presupposes our bodiliness in a quite abstract sense. This is to say, that we experience the body as both *transparent* and *elusive*:

‘We are operating here with a very subtle form of bodily consciousness. Indeed, the body here is not taken as an intentional object of experience, and it is not even experienced in the background or periphery of one’s attentional field.’ (Dorothee Legrand 2013, 215–216)

However, according to Legrand, such experience also involves an awareness of one’s *voluminosity* (roughly aligned with Husserl’s notion of *extension*). It is in this regard that she argues that where Zahavi talks of the ‘zero point’ or ‘here’ of one’s oriented perspective, we ought rather to talk of an ‘orienting volume’:

‘This volume may not be experienced as ‘this specific volume’, and voluminosity may not be experientially contrasted with ‘thinness’ or ‘flatness’. Rather, what matters at this level is that one experiences oneself as a volume *by experiencing one’s location and orientation in space*: through his vision of an object in front of him, the seeing subject experiences himself as having a back and a front, a right and a left, a down and up. In seeing in front of you, you do experience the very fact that you do not see behind your head. In other terms, the seeing subject experiences the very fact that he has a voluminous, localized, and oriented body structured in such a way that only forward-looking is possible.’ (Dorothee Legrand 2013, 217)

(2) When I reach out towards the rose bodily sensations begin to form part of bodily self-awareness, but not in a sense that presents the body as an object. Rather, for Legrand, these sensations provide one with ‘a sense of the opacity of the body’ (2013, 219) (as opposed to the transparency associated with (1)). One is still at the level of the subjective body, but a body not just understood as a volume or orientation, but also, for instance, as solid and receptive to touch. If I am aware of a spider

crawling across my hand the intentional object of my experience will be the spider, but it is an experience uniquely characterised as for myself as a subject receptive to touch, and as touched at a particular location on my body. As such, the unique sensation of the spider adds definition to the contours of my body, even when those contours are not what I am attending to. Similarly, when I reach out for the rose, my intentional object is still the rose, but my pre-reflective experience of self as a subject with the power to move and reach (and as having the bodily apparatus necessary to move and reach with) fundamentally structures that experience. Here, therefore, we begin to see those features enter the picture that were discussed above in terms of embodiment; those forms of familiarity with our bodies that are only open to us. However, they are still not taken as object in their own right, but felt at the level of the pre-reflective character of experiences that are themselves still essentially about a world beyond the body.

(3) When I pay attention to my hand, and indeed when I feel a pain in my finger, I now take my body as my intentional object, but I do so in a way that acknowledges my body as a lived subjecthood. I attend to my hand when I move it, but I attend to it as an object under my control. I attend to my finger when it hurts, but I attend to it in terms of its painfulness for me. In other words, we have ways in which we can take the body as intentional object, but specifically as a way to access the subjective body at a more nuanced level. Legrand and Ravn (2009) illustrated this point perfectly with an empirical study of dancers and the relation they enjoy to their bodies. The point of this study was, in part, to demonstrate that dancers, especially while practicing, pay close attention to their bodies via various perceptual means. The means in question was actually seen to differ between types of dance. For instance, visual perception was most important for ballet dancers (such as by practicing in front of a mirror) so that they could focus on the visual lines of their body, while proprioceptive forms of awareness were more important for contemporary dancers who are more focused on 'how a flow of movement is generated and continued from the way the dancer's specific body functions.' (Legrand and Ravn 2009, 398). But regardless of the specific means, the purpose in each case is to allow the dancers to enhance their pre-reflective understanding of their bodies, that is, to be able to improve as dancers and exert control over how their body will be viewed from a third-person perspective. Thus, their experience takes the form of an attending to their exteriority, but in a way that takes that exteriority as intrinsically tied to their bodily subjecthood.

(4) Finally, when I scrutinize my finger for a hidden thorn, Legrand argues that I take my body as an object in the most extreme sense:

‘Indeed, by taking my finger as an object of *scrutiny*, I do *not* specifically access it in its subjectivity. On the contrary, I experience it in its objectivity: as it is penetrable and penetrated by another object (the thorn), it deploys its belongingness to the physical realm, and I could adopt the same scrutinizing observation to find a thorn in *your* finger.’ (2013, 214)

It is in this sense that Legrand refers to this form of self-awareness as a reification of subjecthood (2013, 224), as in scrutinizing my body I go beyond merely attending to it and render it a thing in the world. In other words, I take a quasi-third-person perspective on it (this particular interpretation will be partly qualified below).

Through this, quite accurate, four-part story Legrand recognises that our bodily self-awareness does not just come in two flavours, as it were, but ranges from the more subjective to the more objective with a host of variations in between, and where a non-ascriptive sense of self as a minimally bodily subject is just one extreme of this range⁶². However, it is not entirely clear how Legrand is to account for this fact. She recognises that our experiences of the body are ambiguous, in the sense that there is always intermingling of subjective with objective aspects, but what is the source of this ambiguity? That Legrand struggles with an answer to this question is shown by the fact that her account demonstrates a persistent tendency towards continued categorisation. This is to say, that she does not fully embrace the idea of the body as a spectrum; she acknowledges the different forms of bodily self-awareness, but still brackets them into either bodily awareness of the self-as-subject ((1) and (2)) or the self-as-object ((3) and (4)). As such, the two categories with which we began still persist to an extent and, I believe, perform some important work behind the scenes of Legrand’s account. By contrast, there are good reasons to believe that the above considerations support the idea that this bracketing ought to be abandoned entirely.

Indeed, what Legrand’s portrayal of the variety of bodily self-awareness shows is the need to fundamentally re-evaluate the distinction between reflective and pre-reflective aspects of experience. For, as has been argued, it is this distinction that cleaves apart subjective from objective forms of bodily self-awareness. On the one hand, one can be aware of oneself as intentional object and in doing so take a reflective stance on one’s being. On the other hand, such an experience, indeed any experience, entails a pre-reflective awareness of oneself as the subject of that experience. As such, if our experiences are suggestive of a less than clear cut division between subjective and objective self-awareness (as Legrand’s account shows) this must amount to a breakdown in the underlying division between the reflective and the pre-reflective. This is not to say that the idea

⁶² For a compelling alternative way in which this variation can be explicated see Legrand (2010), which examines the effects of anorexia on forms of bodily self-awareness.

that there is any such distinction to be made ought to be rejected. Such an approach would thoroughly undercut the position which has been so carefully constructed and amount to throwing the baby out with the bath water. It is instead to claim that pre-reflective aspects of experience must be understood as one pole of a spectrum in which the reflective is the other, and that experience is not to be thought of as composed of two wholly distinct parts, but a range of indistinct and variously reflective elements.

It is the work of Merleau-Ponty, who incorporates elements of *Gestalt psychology* into his phenomenological account of perceptual experience, which provides us with just such an holistic view⁶³. What Merleau-Ponty's account most directly inherits from the former is, firstly, the idea that distinct objects do not appear in our perceptual field in virtue of an association of the various stimuli out of which they are formed. For instance, in perceiving the house across the street out of my window, it is not that I perceive yellow bricks, red tiles, white sash windows etc. and construct out of these visual parts a house. Rather, there is a sense in which the whole must be understood to precede the parts; what I immediately grasp in perception is just a house, with the parts only appearing to me as parts of this pre-given whole (e.g. as the red bricks of that house over there). Secondly, the perception of this whole is only available in virtue of its relation to a background. In other words, my ability to grasp the house as a whole, prior to any recognition of the association between its visual parts, is that this appearance is conditioned by the perceptual field in which it appears and against which it stands out. Those features of the field that are excluded from the figure attended to are as central to the perception of this particular object, as those which are included. The figure, or intentional object, thus both precedes the perception of its parts and is co-determined by the background or horizon by which it is brought into relief:

'If I am walking on a beach toward a boat that has run aground, and if the funnel or the mast merges with the forest that borders the dune, then there will be a moment in which these details suddenly reunite with the boat and become welded to it. As I approached, I did not perceive the resemblances or the proximities that were, in the end, about to reunite with the superstructure of the ship in an unbroken picture. I merely felt that the appearance of the object was about to change, that something was imminent in this tension, as the storm is imminent in the clouds.' (Merleau-Ponty 2013, 40)

⁶³ The extent to which the following is a faithful interpretation is open to debate. For a start, it is based mostly on the position Merleau-Ponty outlines in his *Phenomenology of Perception* (2013) and the question of whether or not he gave up on some of these ideas in his later work *The Visible and the Invisible* (1969) will not be considered at all (cf. Carman (2008, 120–132)). Secondly, it is an interpretation that is guided to a large extent by the ideas that have already been developed above regarding the nature of subjecthood and it is not altogether clear whether or not Merleau-Ponty would have accepted the Zahavian view of things (in large part because his own views on subjecthood are less than clear). Nevertheless, it is faithful to the spirit of charitable interpretation that Merleau-Ponty himself regularly employed and ought to be assessed primarily on the basis of whether it presents a coherent and compelling position.

However, it is clear that for Merleau-Ponty the significance of Gestalt psychology goes beyond an elucidation of the interplay between figure and background in perception:

‘[If] we want to give an unprejudiced definition of gestalt psychology’s philosophical meaning, we would have to say that, by revealing “structure” or “form” as irreducible elements of being, it has again put into question the classical alternative between “existence as thing” and “existence as consciousness,” has established a communication between and a mixture of, objective and subjective, and has conceived of psychological knowledge in a new way, no longer as an attempt to break down these typical ensembles but rather an effort to embrace them and to understand them by reliving them.’ (Merleau-Ponty 1964, 86)

In other words, for Merleau-Ponty, what Gestalt psychology calls into question is precisely the idea that there is any sharp boundary to be drawn between subjective and objective forms of consciousness. Or, more specifically, between reflective and pre-reflective aspects of experience. Experience is, rather, exactly the kind of mixture of reflective and pre-reflective that is required by Legrand’s view of bodily self-awareness as a spectrum.

How is this leap possible? Why does Merleau-Ponty think that a theory about the relation between figure and background in experience effectively breaks down the barriers between reflective and pre-reflective forms of awareness? In many ways, Merleau-Ponty is reading into Gestalt psychology that which he himself adds to it. Thus, to see why he places such significance on it, one needs to understand how Merleau-Ponty views this figure-background structure as itself conditioned through the prism of the body.

For Merleau-Ponty, to be a bodily subject is for one’s experiences to be presented as shaped by a certain *body schema*. In other words, it is this schema that establishes us as a being both in and toward the world. As should already be clear from the previous section, this schema is not to be understood as a definite or concrete awareness of our bodies as an object amongst other objects. Rather, it is closer to that sense in which the way we both interpret and interact with the world is guided by our understanding of the, often vague, contours of our bodily being; an understanding that can only be understood as prior to reflective thought. As is typical of Merleau-Ponty’s work, he sees the importance of this body schema as coming to the fore most clearly in cases of bodily damage or pathology:

‘When an insect instinctively substitutes a healthy leg for the leg that has been removed, it is not, as we have seen, that a pre-established safety mechanism is automatically triggered and substituted for the circuit that has just been put out of service. But no more is it because the animal is conscious of a goal to attain and uses its limbs as different means, for then the substitution would have to be produced each time that the action is blocked and we know that it does not occur if the leg is merely tied. The animal simply continues to exist in the same world and carries itself toward this world with all of its powers. The tied limb is not replaced by the free one because the tied one continues to count in the animal’s being and

because the impulse of activity that goes toward the world still passes through that limb.’
(Merleau-Ponty 2013, 106)

This is to say, that it is not that the insect, upon the loss of a limb, has to rectify their explicit or reflective awareness of their body in order to accommodate the loss. Indeed, it is questionable that they possess such a form of awareness. Nor does it need re-calculate how to use this, now partially unfamiliar, body in order to attain its goals. Rather, it moves towards its goals with those bodily forces that it pre-reflectively understands as available to it. If any re-calculation occurs, this is not of a reflective or conscious sort, but of the sort that is worked out as the animal goes and through processes of trial and error; as it stumbles and falters and learns to cope with the new contours of its form in its attempt to achieve its goals.

To utilise a human example with which he also has a keen interest, Merleau-Ponty says of the patient suffering from a phantom limb:

‘[I]f he treats it in practice as a real limb, this is because, like the normal subject, he has no need of a clear and articulated perception of his body in order to begin moving. It is enough that his body is “available” as an indivisible power and that the phantom leg is sensed as vaguely implicated in it.’ (Merleau-Ponty 2013, 110)

To put it differently, for Merleau-Ponty, a phantom limb arises in those cases when the body schema fails to adapt to a change in the way that occurs so instinctively for the insect. The old contours of the patient’s bodily being stubbornly remain and, insofar as they do so, the present feeling of the limb as a possible force toward the world remains as well. For, being in the world is not just to be passively receptive to that world. Rather, it requires a pre-reflective scheme of things; that physical way in which I take myself as an active bodily force toward the world. Moreover, it is a scheme of things that I take for granted to such an extent that it only reveals itself to me explicitly when it goes awry, that is, when my own pre-reflective sense of self is found wanting by my experiences of the world.

Part of what it means to say that one has such a schema, is that the spatiality of the body is unlike *objective* space. The parts of my body are not a series of points in space and my limbs are not laid out side by side, but already given as enveloped in a body schema. As Merleau-Ponty puts it, ‘the spatiality of the body must descend from the whole to the parts, my left hand and its position must be implicated in an overall bodily plan and must have their origin there.’ (2013, 128) Or, as one might put it, the significance of where my hand is, or at least its significance for me, has nothing to do with its position in objective space. I do not *find* my own left hand four feet from the floor, or even three inches from my right hand. Rather, I *find* my hand in the way it figures in a certain posture I assume or in the way in which I use my body as a force thrown toward a certain task.

When I want to open my office window by means of turning the handle, I do not need to first locate my hand in objective space and then push it toward the handle, as it were. I already *have* my hand in a more immediate and pre-reflective sense, as a part of my body schema understood as an ensemble of powers to act in the world. If there is any sense in which my hand is located when I go to turn the handle, it is the sense in which it is immediately brought to the fore in my using this ensemble:

‘Psychologists often say that the body schema is *dynamic*. Reduced to a precise sense, this term means that my body appears to me as a posture toward a certain task, actual or possible. And in fact my body’s spatiality is not, like the spatiality of external objects or of “spatial sensations,” a *positional spatiality*; rather, it is a *situational spatiality*. If I stand in front of my desk and lean on it with both hands, only my hands are accentuated and my whole body trails behind them like a comet’s tail. I am not unaware of the location of my shoulders or my waist; rather, this awareness is enveloped in my awareness of my hands and my entire stance is read, so to speak, in how my hands lean upon the desk.’ (Merleau-Ponty 2013, 129)

However, most importantly, this is not to say that the bodily subject or its schema is not spatial. It is not that spatiality in the traditional objective sense is to be understood as *real* space, where the spatiality of the body is just something entirely different. Instead, Merleau-Ponty’s claim is that bodily spatiality is, in some sense, primary or primordial spatiality; the spatiality of the body is that by which there is space in the first place and gives rise to the idea of objective space in a sense that is open to reflective thought. As he puts it, ‘our body is not primarily in space, but is rather of space.’ (Merleau-Ponty 2013, 184)

It is at this point that we reach what is perhaps the crux of Merleau-Ponty’s account of perceptual experience. For, to say that the body is *of space* is already to radically alter how we conceive of that which is *in space*, namely, the world of which we have perceptual or intentional experience:

‘Bodily space can be distinguished from external space and it can envelop its parts rather than laying them out side by side because it is the darkness of the theatre required for the clarity of the performance, the foundation of sleep or the vague reserve of power against which the gesture and its goal stand out, and the zone of non-being *in front of which* precise beings, figures, and points can appear. If my body can ultimately be a “form,” and if there can be, in front of it, privileged figures against indifferent backgrounds, this is insofar as my body is polarized by its tasks, insofar as it *exists toward* them, insofar as it coils up upon itself in order to reach its goal, and the “body schema” is, in the end, a manner of expressing that my body is in and toward the world. With regard to spatiality, which is our present concern, one’s body is always the implied third term in the figure—background structure, and each figure appears perspectively against the double horizon of external space and bodily space.’ (Merleau-Ponty 2013, 129–130)

This beautiful passage perfectly summarises the significance that Merleau-Ponty places on Gestalt psychology as that which breaks down the boundaries between pre-reflective and reflective forms of awareness and, thereby, between the subject and the object. For, once the role of the body in

the figure-background structure is understood, one also understands the way in which the space of the world and the space of the body are intrinsically tied up with one another.

To extend upon Merleau-Ponty's own analogy above, one can think of a theatre performance, or at least our ability to perceive that performance properly, as resting on three components. Firstly, you have the actors on the stage (in perceptual terms, the figure). They are what is directly attended to or actively noticed and, in a sense, constitute the substance of the performance. Next you have the stage (in perceptual terms, the background). Mostly we do not need to directly attend to the stage itself, but it is the setting in which the performance of the actors makes sense. If the stage were to be different so would the performance, at least in the respect that the meaning that we would attribute to it would be altered. Lastly you have what Merleau-Ponty discusses as the *darkness of the theatre*, but if we wished to be pedantic this would include aspects such as the silence of the theatre and the direction of the audience (in perceptual terms the body schema). In many ways this darkness is what is never attended to, for there is a sense in which it represents an absence which is reticent to attention. To try to attend to such darkness is not to apprehend it in the same sense in which it structures our perception of the performance, for its role is precisely to draw our attention to that performance and not to anything else. In other words, it would be far too narrow to understand the performance as merely the actors on the stage. Rather, the performance arises out of the totality of these three features and cannot be understood as reducible to any one of them. Perceptual experience is not just the figure or object attended to, nor that figure as part of a perceptual field (i.e. placed within a background). Experience is the whole, and it can only be made sense of as a mixture of figure and background, as well as the body schema through which they are organised:

'One's own body is in the world just as the heart is in the organism: it continuously breathes life into the visible spectacle, animates it and nourishes it from within, and forms a system with it.' (Merleau-Ponty 2013, 245)

'The thing and the world are given with the parts of my body, not through a "natural geometry," but in a living connection comparable, or rather identical, to the living connection that exists among the parts of my body itself.

External perception and perception of one's own body vary together because they are two sides of a single act.' (Merleau-Ponty 2013, 247)

Moreover, it would be a mistake to understand the figure, background, and body schema of experience as distinct parts, the sum of which is a perceptual experience. Part of the thought contained in the short passages above is that the boundaries between each cannot be understood as clearly defined, but as naturally lapsing into one another. To see why, it is useful to consider the case, discussed by Brian O'Shaughnessy (1998), of playing a tennis stroke. In the context in which

this example arises O'Shaughnessy is concerned with whether or not proprioception is a form of perception. To put it differently, the question is whether proprioception is a means of perceiving, for example, the relative positioning of our limbs, or whether it is some more immediate, or perhaps non-conscious, process that provides us with this knowledge. For, as he puts it, what is at least very clear about proprioception is that it is highly 'attentively recessive' (1998, 175) and this makes it difficult to classify in this respect.

O'Shaughnessy argues that there are atypical uses of proprioception, in which one is completely focused on some part of one's body, which are quite clearly a form of perception. For instance, if I take one of my limbs as intentional object, say by attending closely to the position of my arm without looking at it, it is hard to deny that my awareness of its position is of a perceptual variety. However, as he points out, far more typically proprioception is involved in providing us with an awareness of things like limb position when we are not directly attending to our bodies, but absorbed in some form of familiar instrumental action that directs our attention out into the world. And it is these uses of proprioception that are far more problematic in terms of a theory of perception:

'A difficulty arises at once... It is a problem posed by the fact that we have at any instant only so much attention to go around and no more. If perception consists in items coming to the attention, and if proprioception is a variety of perception, proprioception must one assumes make demands upon the attention. But if attention needs in part to be absorbed in proprioception of the acting limb, it rather looks as if when we engage in intentional manipulative action the phenomenon of proprioception ought to be a discordant and distracting item, competing for our attention with both the act itself and (say) visual perception. Yet this is not our experience. When we play a stroke in tennis we are not conscious of a conflict within the attention, we do not experience the limb as competing for our attention with the ball (whose path occupies so much of our attention), nor with the playing of the stroke (which does the same).' (O'Shaughnessy 1998, 178)

To put it simply, the problem is that perception requires attention, but if proprioception takes up attention then it seems as if it is actually going to hinder, rather than aid, the completion of everyday actions, such as playing a tennis stroke, in which it is involved. However, O'Shaughnessy's solution to this problem is actually rather simple. It is, in essence, to point out that attention need not be conceived of in absolute terms, where something is either attended to or it is not, and that the attention required for proprioceptive awareness of our limbs in playing a tennis stroke might be quite minimal:

'If playing a sudden snap forehand volley is intentional not merely under 'hitting the ball to the corner' and 'playing a forehand volley' and 'swinging the racquet' but also under 'moving my right hand forward thuswise', then proprioception of arm position is a necessary condition of such a deed. Even though one's attention is focused primarily on the path of the ball, and doubtless also though to a much lesser extent on the path followed by the racquet,

some small measure of attention must be left over for the movement of the arm.’ (1998, 179)

For O’Shaughnessy, even with such familiar instrumental actions, where the focus of one’s attention extends beyond one’s body, one must still understand proprioception as providing perceptual awareness of our bodies. It is just that the attention taken up by such perception might be very small and perhaps must be understood in terms of, what we might call, a hierarchy of perceptual schemata.

Now, if one takes O’Shaughnessy’s considerations regarding *attention* as roughly synonymous with *reflection*, as seems reasonable, then there is a great deal about his analysis that should appear misguided in light of Merleau-Ponty’s account of perceptual experience. In particular, he does not consider that there might be pre-reflective forms of bodily self-awareness (or what we might call, in O’Shaughnessy’s terminology, pre-attentive, but not un-attentive, forms of bodily self-awareness). It is not a possibility for O’Shaughnessy that the memory I have of my limb movement in playing a tennis stroke may be a consequence of the fact that in taking the flight of the ball as my intentional object and thereby attending to it, I pre-reflectively take my bodily movement as the subject of that attention. For, on the account that sees the subject as intrinsically bodily, there is no reason to think that all the awareness involved in playing a tennis stroke must be an explicit intentional attending. To put it differently, O’Shaughnessy assumes that the only way in which my body is open to perception is in taking it as intentional object, that is, to possess an awareness *of* my body. What Merleau-Ponty demonstrates is that this is already dependent on the way in which I experience *with* my body.

On the other hand, what O’Shaughnessy’s description of the playing of a tennis stroke does recognise is that attention, or reflection, is a matter of degree rather than all or nothing. The choice is not between paying attention to something or not paying attention to something, because, as he quite rightly notes, I can pay more or less attention. Moreover, just because I attend to a high degree to one thing does not mean that I cannot pay attention to some other thing to a lesser extent. When I play a tennis stroke the locus of my attention is more than likely focused on something like the path of the ball (provided that I am anything but a novice tennis player), but this is not all the attention there is to go around. The same must be held to be true of all experience. When I attend to an object this attention is not like a spotlight that casts everything around it in to darkness. Rather, there is a sense in which attention radiates out from this locus with those features closest to the action, as it were, firmly placed in the foreground, while less necessary features recede into the background. Nor is the figure to which I attend a single wholly determinate entity in an ocean of

vague background. Rather, my attention flits from object to object, and parts of the background are continually coming to the fore or receding as I do so. In other words, it is not that I attend to the object or figure and not to the world or background out of which this determinate being appears. For, those objects closest to the locus of my attention will always be brought into slightly sharper focus than those further away.

This, I hold, is the basis on which Merleau-Ponty takes the figure-background-body structure of experience to constitute a breakdown in the dichotomy between pre-reflective and reflective aspects of awareness. It is not just that the intentional object appears against, as he puts it, the *double horizon of external space and bodily space*; it is that these horizons, as well as the object that is structured by them, must be understood as non-discrete component parts of a single experiential spectrum. In other words, everyday experience is not entirely like being in a theatre in the sense discussed above, in which the main elements are relatively fixed. In fact, part of the purpose of the theatre is to focus and maintain our attention in a way that the ordinary world does not. As such, everyday experience is more like being in a theatre in which the levels of lighting, the direction of the seating, the width of the stage, and the actors on the stage, are constantly in flux. In ordinary perceptual experience the background in which objects appear is always changing as we move through the world, along with the objects to which we are attending and the aspects of our bodily schema through which these objects and background are interpreted.

Again, this is not to deny that there is any distinction to be made between reflective and pre-reflective aspects of awareness (or between the figure, background, and body schema of experience). Rather, it is simply to point out that there is a sense in which actual experience occurs between these poles. There is the explicit intentional object to which we are attending at any given moment and there is the most minimal form of bodily self-awareness that is and must be given in every such intentional experience. However, there is also everything in between and in terms of which both of these extremes can be made sense of.

Thus, to return to the point from which we began, it is the fact that experience itself must be understood as just such a spectrum that allows us to account for Legrand's assertion that bodily self-awareness exists on a spectrum from the wholly subjective to the more or less objective. Our reflective awareness has a focal point, but it casts its light far wider than this point. When I take some feature of the world not directly related to myself as my object (such as a rose in the distance), this point is focused far from myself and, as such, my bodily self-awareness recedes into the background; a mere shadow of that which is reflected upon. All that is left is a vague and pre-reflective awareness of a certain *orienting volume* as Legrand put it. When I take a feature closer to

myself as object (such as a spider crawling across my hand) this focal point is still aimed at the world, but it is drawn closer to myself and its radiating light begins to illuminate the hazy outline or contours of my bodily self. When I take some feature of my own body as explicit intentional object (such as the proprioceptive awareness of my lofted arm) the focal point of my attention is directed firmly back on myself and what was a hazy outline becomes a distinct and determinate concrete entity, and it is the rest of the world that recedes into the background (along with those features of my body not being directly attended to). In bringing my bodily self under the gaze of this focal point a certain transformation takes place; the experience of myself as a bodily subject when this subjecthood is thoroughly submerged in the background must differ from my experience of it at the centre of my attention. For, in the second instance I am having an experience *of* my body, whereas in the first I am having an experience *with* my body. But this transformation should not be thought of as a sharp ontological distinction that renders the body under the gaze of reflection part of a completely distinct ontological category from that which does the gazing. Rather it must be thought of as two ends of the same spectrum, for my intentional gaze can be brought further from or closer to myself.

What then, on this account, is to be said regarding the place of pure non-ascriptive subjective self-awareness, that is, the for-me-here-ness of experience? Legrand's account suggests that this is what we are left with when we attend to objects that are, in a sense, farthest removed from our bodies, but this is too simplistic a characterisation of how we experience the world. Legrand's neat and distinct four-part story above in itself must be a kind of caricature of the complex and interwoven way in which experience is presented to us⁶⁴. It is more likely that the way in which we do experience is so intimately tied to our being active moving subjects in the world, that any experience is going to be bodily in a sense that goes beyond a mere for-me-here-ness. However, this is not to suggest that there is no such character to our experiences. Rather, it is to say that all experience is characterised by such a subjective node, upon which our progressively more concrete experiences of the bodily self are formed. The possibility of having an experience the subjective character of which is simply for-me-here-ness might be an abstract possibility, but it is one which acts as the most fundamental way in which our experiences are structured. Experience might never strip the bodily subject down to this point, but it is always there nonetheless and it is that which renders us an individual bodily subject in the first place.

⁶⁴ Legrand herself does appear to be well aware of this: 'This simple scenario is characterized by four different forms of bodily self-consciousness. Notice that some or all of them are paradigmatically lived in conjunction with one another but here I present them successively to better analyse their specific characteristics.' (2013, 213)

What about the opposite end of the spectrum and the possibility of being self-aware of the body purely as a physical object? If a pure for-me-here-ness remains an abstract and unlikely potentiality, this must be more so in the case of encountering oneself purely in terms of one's exteriority. Legrand might describe her example of scrutinising one's finger for a hidden thorn as a taking of the body in the same way that I might encounter the body of the other, but this seems unlikely. It is hard to imagine that I can ever completely remove my understanding of my finger as being a part of me from the equation. After all, scrutinising one's own body never has the same feel as scrutinising someone else's (and even in the latter case, it is significant that scrutinising someone else's body is never the same as scrutinising an inanimate object). Moreover, it is important to note that when I take my body in the very objective sense, it is not that the subjective character of this experience disappears. Rather, my experience must still be presented as for a more or less defined bodily subject and one suspects that this fact, along with the fact that my scrutinised finger will be experienced as continuous with this bodily subject in some sense, will prevent me from ever taking my own finger in the same sense that I might take yours. Perhaps, therefore, the body as a pure exteriority remains a theoretically understood principle that structures the ways in which I experience myself, even if it can never be practically realised in experience.

Regardless, what is crucial is that the body must be understood as that which unites subjective and objective self-awareness. All self-awareness, it has been argued, is necessarily bodily, for even my most minimal and pre-reflective experiences of self are bodily in some sense. However, it is the ambiguous variations in this sense that account for the awareness we have of ourselves both as a minimal and abstract subject and as a concrete being in the world. Moreover, and in virtue of this, we are in a position to provide an answer to the question with which we began this section. The bodily subject can be understood as both a more or less determinate entity and as a minimal non-ascriptive entity, precisely because these do not correspond to two forms of being but different parts of the same bodily spectrum. At a certain level all experience is structured by the minimal and non-ascriptive sense of bodily subjecthood that we have, not as a concrete entity, but simply as a subject in the most abstract sense. However, experience extends out from this pole towards the world of intentional objects, and in doing so the bodily subject finds itself as an active being and an ensemble of powers toward that world, that is, a particular body schema.

IV: Conclusions

The way in which the above argument has been made not only performs crucial work in terms of the direction of this thesis, but also approaches the topic from an angle that is rarely broached in the literature. Indeed, phenomenologically speaking, the place of the body usually arises in the context

of trying to account for our encounter with other subjects (a fact that will be discussed in some detail in Chapter 6). As such, the individuality of the subject is often a presupposition of the discussion, the task of which is to open this individuality up to access, and be accessed by, others. By coming to this discussion with the problem of individuation itself unresolved, one is forced to take a far more radical approach to the topic. Exteriority can no longer be approached as an afterthought, tacked on to a subject that is more fundamentally an interior, and the reflective cannot be denigrated as a form of experience less primordial than the pre-reflective. Instead, having an exterior must be understood as part and parcel of being an interior and thinking reflectively must be understood as only differing in degree from experiencing pre-reflectively. To be a subject, on this view, is to simultaneously be both a subject and object for oneself, as those mutually reinforcing poles that constitute one as both distinct from and toward a world.

Moreover, while it was suggested (in section II) that an understanding of subjective self-awareness as constitutively embedded is sufficient to individuate a particular subject in virtue of its intrinsic spatial locatedness in the world, it should now be recognised that, in many ways, the Merleau-Pontian spectrum view provides us with a far more robust solution to the problem of individuation. On the account that has been sketched, not only is the subject individuated in some minimal way in virtue of the *for-me-here-ness* of their experience, but this most subjective pole of their self-awareness must already be understood as part of the same spectrum as more concrete forms of self-awareness. The precise features of this spectrum will vary from person to person, and the extent to which one grasps oneself as a concrete physical entity might, for instance, develop in infancy. However, in all but the most abstract cases, to be subjectively self-aware will be to already *have* oneself in less subjective and more objective ways. To put it differently, it was suggested earlier that, while the subject-as-embodied cannot achieve the desired individuation of the subject, it does nevertheless seem to grasp something important about the ways in which we do experience our bodies. What the spectrum view allows us to do is to link these insights with ideas regarding the more primordial role of the body in self-awareness that are captured by the subject-as-bodily view. The result is an account that is both able to make sense of the subject as bodily, and individuated at even the most basic level in virtue of that bodiliness, as well as link this basic sense of bodiliness on a seamless spectrum with more concrete and reflective forms of bodily self-awareness.

With this individuation of the subject established, the following chapter will turn to the second part of our dual task. The question is whether or not this account of individual subjecthood leaves sufficient room such that it might be understood as capable of being shared and, therefore, open to pluralisation. Or has the above account, in successfully individuating the subject through its

bodiliness, effectively closed off any such avenue of thought? For, there might be a strong intuition to the effect that if to be a subject is to be a bodily entity, then this is not going to be a form of subjecthood open to groups.

Chapter 6

The Plural Bodily Subject

What has been shown, thus far, is that subjective self-awareness must be understood as individuating in virtue of its intrinsic bodiliness and that, in turn, this bodiliness places self-awareness, and the subject that it constitutes, on a spectrum from the most subjective to the most objective. The question that must now be confronted is whether the subject so constituted is open to pluralisation? Are there plural as well as individual subjects?

In Chapter 3 an argument was considered, put forward by Hans Bernhard Schmid, to the effect that in the same way that individuals are constituted by experience with a first-personal subjective character, plural subjects can be thought of as constituted by experiences with a first-personal-plural subjective character. The problem that was attributed to this argument is that it rests on the assumption that such a subjective character is a form of self-referential attitude; a mistake that, in turn, uncovered the problem of how such a subjective character is capable of constituting a *particular* subject in either the individual or plural case. However, in coming to grips with the idea that individuation is realised by the intrinsic bodiliness of subjective self-awareness, are we now in a position to defend the idea at the core of Schmid's position, namely, that there is experience with such a subjective character that it can be understood as constitutive of a plural subjecthood?

On the contrary, rather than getting us closer to an account of plural subjective self-awareness, the idea of such awareness as bodily appears, at least on the face of it, to make such an account even more problematic. If to be a subject is for one's experiences to be given from a certain bodily perspective, it is hard to envision how such a perspective can be thought of as being possessed by a group. Whatever qualities groups might be thought to possess, intuition tells us that a single unified body, on which such a unified perspective might be founded, is not among them.

Nevertheless, the aim of this chapter is to break down precisely this intuition and provide a speculative account of how groups might be understood as being continually in the process of developing just such a unified bodily experiential perspective. This is to say, it will be argued that while individuals can be thought of as developing in childhood, and indeed throughout their lives, both an awareness of their experiential world and themselves as the bodily subject of that world, groups can be thought of as developing in something like the opposite direction: from a concrete awareness of themselves as a distributed collection of individual bodily subjects, towards an

awareness of themselves as the single unified bodily subject of that concrete existence. To be a plural subject, on this view, is to be a group that is permanently progressing toward individuality.

Moreover, it will be suggested that while Merleau-Ponty had nothing to say directly on the possibility of plural subjects so constituted, he nevertheless provides us with most, if not all, of the theoretical tools necessary to defend just such a possibility. Specifically, while Merleau-Ponty's concern, like most phenomenologists, was with intersubjectivity as opposed to plural subjecthood, his explanation of the conditions on which the former occurs provides us with the kind of anti-solipsistic view of subjecthood that is necessary to account for the possible *sharing* of subjecthood. To this end it will prove necessary to first consider Merleau-Ponty's concept of the habit body (section I), which grounds the idea that, if the subject is to be understood as intrinsically bodily, then it cannot be understood as confined to the skin and bones of the individual organism. In other words, to be a bodily subject, on this view, is to be naturally extended out into the world. It will then be suggested (section II) that in combination with his understanding of the cultural world, this view of the body constitutes the basis of Merleau-Ponty's account of intersubjectivity and his response to the problem of other minds; the existence of others does not pose a problem for us because we both presuppose the existence of others and are naturally extended out into a world of shared bodily spaces. While, as an account of intersubjectivity, it must be recognised that this in itself does not constitute a defence of plural subjects, and ultimately Merleau-Ponty's own thinking remains locked within the categories of self and other (section III), it will nonetheless be argued that the ideas, contained therein, provide us with the basis on which to defend the thought that we can understand groups as continually developing towards unified subjective self-awareness through processes of habitual familiarisation (sections IV and V).

It is in this final claim that the thrust of this entire thesis rests. We began with the question of whether groups of individuals are ever capable of thinking or acting in ontologically collective ways. The aim of this chapter is to definitively answer this question in the affirmative. Groups of individuals can both think and act in ontologically collective ways in virtue of the fact that it is at least possible that they can be subjects in the same way that those individuals themselves are subjects; they can have experiences in which they are presented as the pre-reflective bodily subject of that experience. Once we understand Merleau-Ponty's idea that subjective self-awareness is wrapped up in the idea of having a certain body schema, and once we understand such a schema as being forged through process of habituation, the idea that groups can also go through such processes of habituation and, thereby, become phenomenally self-aware, ceases to appear quite so counter-intuitive. However, showing that this is the case will involve tempering this claim in two respects:

(1) plural subjects must be understood as developing in something like the opposite direction to individual subjects, that is, from reflective to pre-reflective self-awareness; (2) plural subjects can never achieve the kind of subjective self-awareness enjoyed by individuals, because to do so would be to become an individual, and it is not clear that *this* is a kind of subjecthood that is ever enjoyed by groups.

I: Merleau-Ponty on Habit and Extension

The first thing that must be recognised is that Merleau-Ponty's explication of the bodily nature of experience, and the subject of experience, is not intended to produce, what we might call, a *bodily solipsism*. This is to say, that what Merleau-Ponty is concerned to present is a view of the subject as naturally open to, and extended out into, the world in which they are constitutively embedded.

To see why let us recall that in the last chapter Merleau-Ponty's account of the bodily subject was described as being based around his notion of the body schema and its unique spatiality. For Merleau-Ponty, to be a bodily subject is for experience to present itself as for a being both in and toward the world, and to be presented as toward a world is to grasp and confront that world through the prism of a certain body schema. To recall the example of the ant which has a leg removed, it confronts the world and goes about its projects in terms of the affordances provided by the body that it has to hand; affordances that change when the leg is removed and the ant is forced to compensate accordingly.

What has not been touched on thus far, however, is that, for Merleau-Ponty, this idea of a body schema is tightly bound up with the concept of *habit*. In fact, Taylor Carman goes as far as to suggest that the terms 'schema' and 'habit' can be taken to be virtually co-referential for Merleau-Ponty (Carman 2008, 106). Thus, the sense in which the ant has its body *to hand* is the sense in which it is habitually familiar with its body in a way that enables action without reflective thought, and it is re-habitation that allows it to take the loss of its leg in its stride (an habituation that, in the case of the ant at least, merely requires a few preliminary errors before the body is rediscovered, as it were). Similarly, it is habitual familiarity that allows us to complete a plethora of mundane movements without the slightest extraneous thought:

'If I am told to touch my ear or my knee, I bring my hand to my ear or to my knee by the shortest path without having to imagine the position of my hand at the outset, the position of my ear, or the trajectory from one to the other ... When I bring my hand toward my knee, I experience at each moment of the movement the realization of an intention that did not aim at my knee as an idea, or even as an object, but rather as a present and real part of my living body, and ultimately as a point of passage in my perceptual movement toward the world.'
(Merleau-Ponty 2013, 180)

This unique familiarity with our own bodies should not come as a surprise. After all, what Merleau-Ponty's argument is intended to show is that *I am my body* in the most important sense. Nevertheless, with regard to present purposes, it is crucial to recognise that what I am, therefore, is something that is capable of undergoing change. Of course, there is also a sense in which at all times an immutable abstract subjecthood must remain, that is, the for-me-here-ness of all experience. For, this is what it is to be a subject at the most minimal level. However, everything beyond this most extreme subjective pole of the spectrum of bodily self-awareness is at least potentially open to variation. In some particularly extreme cases this variation can manifest itself in certain pathologies. For instance, while the ant, urged on by its environment and the pressing demands of existence, is forced to take the loss of its leg in its stride, for humans such a loss inevitably has a far more dramatic impact. Thus, as outlined above, Merleau-Ponty describes the phenomenon of phantom limb syndrome as a breakdown or lag in the body schema. But our body schemas also vary in far more ordinary ways on an everyday basis, such as in the acquisition of a new skill (or the loss of an old one) or by gaining strength or agility through regular exercise. As such, when I reach with my hand to touch my ear, perhaps to rub my earlobe, it is not that the potential of such action is demanded by the mere fact of experience, for a bodily subject does not have to be constituted much different from myself in order to not have this potential. This is to say, that the particular *shape* of my body schema is not a brute fact of experience, because it could always have been otherwise and because it is always open to change.

The reason why such action does not require reflective thought, and the reason why I have the particular more or less determinate body schema that I do have, is because I am habitually familiar with it. To see the force of this, one only has to consider the disconcerting effect it would have to reach with one's hand to rub one's earlobe and not to find that earlobe, perhaps not to find an ear at all; an effect that would surely go beyond the feeling that one has undergone a serious injury. For, I do not *look* for my earlobe with my hand. Rather, as a bodily subject with a certain body schema, I am this earlobe and this hand, and to reach with my hand for my earlobe and not to succeed is, in a sense, to have unexpectedly lost myself. In other words, such an experience would be so disconcerting because I do not have to find my body at each moment anew, but have my body to hand at every moment because I am habituated to it and to the way that it structures how I exist toward the world. Thus, habituation is that process, whatever it might consist of, that allows us to *have* our own bodies and to become familiar with the contours of our own bodily being in a way that structures our experience of the world; contours that are only brought under scrutiny when

our experiences of the world present themselves as for a bodily subject that differs from that with which we were acquainted.

Furthermore, and more importantly from the perspective of present purposes, it is part and parcel of this understanding of the body schema as habituated that there is no reason to view the boundaries of that schema as aligned with the edges of one's skin and bones. For, after all, the lived or subjective body is not to be understood as a determinate physical object and it should not, therefore, be thought of as continuous with, for instance, a particular individual organism. Rather, the body schema is just that which plays the role of the *implicit third term* in the figure-background structure of experience and, as we have seen, on Merleau-Ponty's account there are no clear divisions to be drawn between these three constitutive aspects of experience. In light of this, to be a bodily subject is to be naturally extended out into the world.

Merleau-Ponty offers the example of the blind man's cane as a particularly stark example of how the body schema can incorporate aspects of the world in this way. For, he suggests, the blind man's cane is not an object among other objects for him, which he uses instrumentally to infer details about the world. Rather, it has become part of his own particular sensory apparatus; as fundamental a part of the way in which he perceives as eyes are to those with sight:

'The blind man's can has ceased to be an object for him, it is no longer perceived for itself; rather, the cane's furthest point is transformed into a sensitive zone, it increases the scope and the radius of the act of touching and has become analogous to a gaze. In the exploration of objects, the length of the cane does not explicitly intervene nor act as a middle term: the blind man knows its length by the position of objects, rather than the position of objects through the cane's length. The position of objects is given immediately by the scope of the gesture that reaches them and in which, beyond the potential extension of the arm, the radius of action of the cane is included.' (2013, 178–179)

In a far more mundane way, our everyday lives are permeated with the constant use of tools and artefacts with which we are habitually familiar, but which do not need to form such an intimate part of our lives as the blind man's cane:

'Without any explicit calculation, a woman maintains a safe distance between the feather in her hat and objects that might damage it; she senses where the feather is, just as we sense where our hand is. If I possess the habit of driving a car, then I enter into a lane and see that "I can pass" without comparing the width of the lane to that of the fender, just as I go through a door without comparing the width of the door to that of my body. The hat and the automobile have ceased to be objects whose size and volume would be determined through comparison with other objects. They have become voluminous powers and the necessity of a certain free space. Correlatively, the subway door and the road have become restrictive powers and immediately appear as passable or impassable for my body and its appendage.' (Merleau-Ponty 2013, 178)

So, to say that the body schema is formed through habit is in part to return to the idea of the unique spatiality of the body. To habituate oneself to one's arms or legs, to a cane, or to a new car, is in each case to subsume that thing within the spatiality of the body. To put it differently, at a certain point in the process of habituation, the thing in question ceases to be an object *in space* and becomes, in joining with the unique spatiality of the body, *of space*:

'To habituate oneself to a hat, an automobile, or a cane is to take up residence in them, or inversely, to make them participate within the voluminosity of one's own body. Habit expresses the power we have of dilating our being in the world, or of altering our existence through incorporating new instruments.' (Merleau-Ponty 2013, 178–179)

'When the typist executes the necessary movement on the keyboard, these movements are guided by an intention, but this intention does not posit the keys as objective locations. The subject who learns to type literally incorporates the space of the keyboard into his bodily space.' (Merleau-Ponty 2013, 180)

Moreover, the sense in which objects can become incorporated through habituation into the body schema, is in no sense tied to the particularity of the object. It is not, for instance, that the blind man is habituated to one particular cane, which as a physical object comes to form part of him. Rather, he becomes habituated to the use of canes *such as this one*. Of course, he more than likely has a cane of which he is particularly fond, but if that cane was to be substituted with another suitably similar cane, it wouldn't be the case that his ability to navigate through the world would suddenly be lost. Likewise, we are all habituated to the use of such a wide range of everyday instruments, and to such a high degree, that we move through life without ever having to consider the sense in which we are able to take them for granted.

To give the most commonplace of examples, consider the fact that on any given day one opens countless doors by means of some form of doorknob or handle. Many of these will, of course, be quite familiar, but an equal number will not. In either case, it is significant how little they reckon in one's attentive experience of the world. Unless a doorknob is particularly alien to me (being of a construction that I have never seen before) I rarely, if ever, take any explicit notice of them. I need not stop to reflect upon their role in my opening of the door and if such reflection were required such routine tasks would take on an unnecessary level of complexity. Rather, the door presents itself to me through my intention of opening it and if I am attentive to anything it will be this goal. Thus, I move toward and through the door without needing to pay attention to the doorknob, without which the door certainly could not be opened. Much like O'Shaughnessy's example of the tennis stroke considered in Chapter 5, it is not that I turn or push the doorknob unintentionally. Indeed, it must be said that I move purposefully to open the door by means of this doorknob and it is more than likely that I will be able to recall my intentionally turning the doorknob

after the event. However, this is not because I had a reflective or attentive experience of using the doorknob, but because it, as an instrument of the most basic habitual familiarity to me, is automatically incorporated into the spatiality of my body as I move through the world. It seems absurd to say, but I need not consider doorknobs in the same sense in which I need not consider the position of my arm when I open the door. For, there is some sense in which the doorknob forms part of my body schema at that moment; just one of the countless instruments that pulse through me as I move toward a world beyond them. Perhaps, however, Merleau-Ponty's own example on this point is even more illuminating:

'It is said that an experienced organist is capable of playing an organ with which he is unfamiliar and that has additional or fewer keyboards, and whose stops are differently arranged than the stops on his customary instrument. He needs but an hour of practice to be ready to execute his program ... He sits on the bench, engages the pedals, and pulls out the stops, he sizes up the instrument with his body, he incorporates its directions and dimensions, and he settles into the organ as one settles into a house. He does not learn positions in objective space for each stop and each pedal, nor does he entrust such positions to "memory." During the rehearsal – just as during the performance – the stops, the pedals, and the keyboards are only presented to him as powers of such and such an emotional or musical value, and their position as those places through which this value appears in the world. Between the musical essence of the piece such as it is indicated in the score and the music that actually resonates around the organ, such a direct relationship is established that the body of the organist and the instrument are nothing other than the place of passage of this relation.'

(2013, 180–181)

It can now be stated, therefore, that to be a bodily subject, and to have a body schema in virtue of habituation, is to be intrinsically extended out into the world. Which is as much as to say that once it is realised that the body is not first and foremost an object, but the most intrinsic ground of our being in and toward a world, it must also be realised that where the body lies, as it were, is not a question of objective space. Rather, it is a question of how that objective space is structured by the spatiality of the body and, in turn, how that spatiality is determined by processes of habituation; it is a question of where our familiarity with the various ways in which we can be toward a world place the, often vague and ambiguous, contours of our bodily being.

II: Merleau-Ponty on Intersubjectivity

So, the bodily subject should not be understood as being limited to the skin and bones of the biological organism. However, Merleau-Ponty makes a claim that is, potentially, even more significant from the perspective of this thesis, in that he argues that not only is the subject intrinsically extended into the world, but that this world in its own right must be understood as immediately social. This is to say, for Merleau-Ponty, to be a bodily subject is to be naturally extended into a world into which other bodily subjects are also extended.

Like most phenomenologists, Merleau-Ponty's understanding of the social world is supposed, in large part, to account for what is known in epistemology as the problem of other minds. To put it in familiar terminology, the question is, if I understand myself as a particular bodily subject, how can I know that there are other bodily subjects like me? Merleau-Ponty's immediate response is, essentially, to deny that there is any such problem in the first place. Specifically, it is to argue that the very fact that such a question presents itself shows that one has already presumed the existence of others and, therefore, negated the supposed force of the problem. After all, the world onto which our perspective opens up, and into which we are extended, is fundamentally a cultural world:

'Just as nature penetrates to the center of my personal life and intertwines with it, behaviors also descend into nature and are deposited there in the form of a cultural world. Not only do I have a physical world and live surrounded by soil, air, and water, I have around me roads, plantations, villages, streets, churches, a bell, utensils, a spoon, a pipe.' (Merleau-Ponty 2013, 405)

It is this appearance of the world as cultural from the offset that presupposes the existence of the other. For, there could be no distinction between a natural world and a cultural world unless I already understood these cultural objects as being tied to a subjecthood other than my own:

'Each of these objects bears as an imprint the mark of the human action it serves. Each one emits an atmosphere of humanity that might be only vaguely determined (when it is a matter of some footprints in the sand), or rather highly determined (if I explore a recently evacuated house from top to bottom).' (Ibid.)

It is, therefore, absurd from the offset to doubt the existence of other minds, for it is intrinsic to the world, as we experience it, that it is experienced as social⁶⁵. Instead, the more important

⁶⁵ It might be argued that, while this is sufficient to dismiss a very broad scepticism regarding other minds, it is not sufficient to dismiss the more particular epistemological problem of other minds, that is, how can I know that the specific person standing in front of me right now is minded (i.e. not a zombie)? Viewed in this way, there are in fact two epistemological problems of other minds and, so far, Merleau-Ponty has only responded to the first. However, it is crucial to acknowledge that Merleau-Ponty does appear to respond to the second problem as well, and his response accounts for a significant part of his theory of intersubjectivity. Specifically, Merleau-Ponty brings in the idea of an anonymous or pre-personal subjecthood; a subjecthood that is prior to individuality and allows us to share a certain pre-personal core with everyone else. It is this shared core, he appears to believe, that allows us to immediately recognise the subjecthood of others.

Nevertheless, this is an aspect of his theory that will be completely avoided here. To those familiar with Merleau-Ponty's work, this might appear particularly flippant, considering the emphasis that is usually placed on this aspect of his view. It is important, therefore, to provide some justification for this lacuna. There are, I suggest, at least four substantial reasons to discard Merleau-Ponty's ideas regarding the anonymous or pre-personal subject, none of which can be defended in sufficient detail here: (1) It is not clear that this view is internally coherent and Merleau-Ponty appears to flagrantly contradict himself on this matter repeatedly; (2) Once we have accepted the Zahavian account of subjective self-awareness, it is clear that the idea of an anonymous subjecthood must be rejected, as it would appear to involve the thought that there can be experience, specifically that of young infants, that is not presented first-personally; (3) As is demonstrated by this section in general, it is perfectly possible to put forward a coherent interpretation of Merleau-Ponty's account of intersubjectivity without any recourse to the idea of anonymous subjecthood; (4) If it is correct that Merleau-Ponty is compelled to assert the idea of such an anonymous subjecthood in order to respond to the second epistemological problem of other minds discussed above, then his motivation is ill-founded because there is no need to respond to this problem. This final point is obviously pertinent to the present

question for Merleau-Ponty is not the epistemological question of how we are capable of knowing that there are other subjects, because this could hardly be doubted, but the conceptual question of how I am capable of grasping the idea of a subject, which is so wrapped up with the individual, as applying to another:

‘How can a human action or thought be grasped in the mode of the “one,” given that it is, in principle, a first person operation and inseparable from an I? The easy response is that the indefinite pronoun is here simply a vague formula for designating a multiplicity of I’s, or even an I in general ... But this is precisely the question: how can the word “I” be made plural? How can we form a general idea of the I? How can I speak of another I than my own? How can I know that there are other I’s? How can consciousness, which as knowledge of itself is, in principle, in the mode of the I, be grasped in the mode of the You, and thereby in the mode of the “One”?’ (Merleau-Ponty 2013, 405–406)⁶⁶

For Merleau-Ponty, the truly difficult feature of other minds to explain is what is really contained in the concept of the *other*. This is, more or less, what is often referred to as the *conceptual problem of other minds*: what is it to have an idea of a subject that is not me, if to be a subject is simply to be presented to oneself as such in experience?

However, as Merleau-Ponty points out, the account of the subject as bodily already provides us with a response to this question. For, that the world is fundamentally encountered in this way is to be explained by the fact that we are not a solipsistic or hermetically sealed internal being (i.e. a mind distinct from a body), but that we are, at even the most basic level, a *bodily* subject. As such, Merleau-Ponty suggests that this problem of having a concept of the other is one which only arises for empiricism and intellectualism, which force the individual to reconcile the irreconcilable when it comes to conceiving of there being another like themselves. This is to say, that if a distinction is made between the physical person and their consciousness, and if one is only, strictly speaking, the latter, then a sort of impossible leap is required in order to recognise others as minded. For, I am only ever presented with the other as a physical presence, and yet what I am to myself is precisely not such a physicality, but an interior consciousness. Thus, in order to attribute a mind or subjecthood to the other I am required to impute to this physical presence, perhaps via an inference

discussion, and amounts to the claim that only the first epistemological problem of other minds is a genuine issue. The second, on the other hand, asks for the removal of doubt regarding something about which we cannot possibly be certain. This is to say, it will always be a possibility that the person standing in front of me is in fact a zombie, and a theory that attempts to eradicate that possibility entirely must be treated as misguided. In this regard, Merleau-Ponty’s argument that it cannot be doubted that there are other minds in general, deals with the only substantial epistemological worry. If, therefore, Merleau-Ponty introduces the idea of an anonymous subjecthood in order to deal with the second problem, and there is good reason to think that he does, then his mistake was in being concerned in the first place.

⁶⁶ Merleau-Ponty’s formulation of the question as, ‘How can I know that there are other I’s?’ in the middle of this passage, should be read as an error. It is quite clear in the context of the discussion that Merleau-Ponty is dealing here specifically with the conceptual problem of other minds, not the epistemological problem (although he does have a tendency to conflate the two).

from signs or gestures, the sort of interior life, which is precisely that which is denied by its mere physicality for me:

‘I would simultaneously have to distinguish him from me, thus placing him in the world of objects, and think of him as conscious, that is, as this type of being without an outside and without parts to which I only have access because I am this consciousness and because he who thinks and he who is thought merge in him.’ (Merleau-Ponty 2013, 407)

By contrast, in rejecting the distinction between the mental and the physical, the subjective and the objective, or what Merleau-Ponty calls the *for-itself* and the *in-itself*, no such *contradictory operation* is required. On his account, the subject is not distinct from its physicality. Rather, subjecthood exists in the spectrum of bodily self-awareness and there is no sharp line to be drawn between the self as subject of experience and the self as a concrete being in the world. Moreover, the bodily subject is intrinsically tied to a world in which it is constantly finding itself; it is open to that world in a way that allows it to extend into it and it is not distinct from the world into which it extends. For Merleau-Ponty, it is this very open-endedness and extension into a world that I must, in virtue of culture, understand as shared with others, which provides me with an innate familiarity with the other:

‘There is, between my consciousness and my body such as I live it, and between this phenomenal body and the other person’s phenomenal body such as I see it from the outside, an internal relation that makes the other person appear as the completion of the system. Others can be evident because I am not transparent for myself, and because my subjectivity draws its body along behind itself.’ (2013, 410)

To put it succinctly, the reason why I can have a concept of a subject other than myself and believe that in encountering another body that I am encountering another subject, is because our subjecthood is not a form of interior life but a form of being that is, as it were, lived out in the open. Thus, to encounter another body is already to encounter it *as a subject*, and I do not have to infer the existence of the other from certain features of an objective world, because *subjecthood is bodiliness*.

It does not take a great deal of insight to understand why these ideas might be significant from the perspective of an account of plural subjective self-awareness. The problem, with which this chapter began, was that if subjective self-awareness must be understood as bodily, then it is hard, at least on the face of it, to envision how such an awareness might be possessed by groups. But Merleau-Ponty’s accounts of habituation and intersubjectivity appear to work against this intuition. For a start, the body is not to be thought of as fixed in terms of a certain concrete physical or objective space. Rather, the spatiality of the body is a more fundamental form of spatiality and the space of the body must be thought of as, through processes of habituation, naturally extending out into the

world. Secondly, this world into which we are extended must be understood as intrinsically social. I have a cultural world just as much as a natural world and I possess an innate familiarity with the other in virtue of our mutual bodiliness. Consequently, the bodily subject can be understood, for Merleau-Ponty, as existing in a world of shared bodily spaces. For, we are continuously drawing cultural tools and instruments into the spatiality of our body, and we understand these objects as bearing the mark or imprint of the other. In this way, we extend into an intersubjective world just as others do likewise and we inhabit and incorporate parts of that same world into our subjecthood. We dilate our being, as Merleau-Ponty might put it, into areas where the other is already dilated, and it is this mutual overlapping that would seem particularly conducive to the idea that we are capable of being party to collective forms of experience with a plural subjective character.

III: The Problem of Self and Other

However, while certainly significant, such an understanding would appear to miss that which is genuinely problematic about Merleau-Ponty's position from the perspective of an account of plural subjective self-awareness. Namely, that his is an explication of *intersubjectivity* and as such it presupposes the sovereignty of the individual subject at some level. This is to say, that as an account of how one subject encounters another, regardless of whether this encounter occurs in shared bodily spaces, it still must assume that what we are dealing with is distinct individual subjects. Thus, no matter how conducive the things Merleau-Ponty says might seemingly be to the idea that subjecthood is pluralisable, this is exactly what is denied by the fact that his is an account of the interaction between self and other:

‘The other’s grief or anger never has precisely the same sense for him and for me. For him, these are lived situations; for me, they are appresented. Or if I can participate in this grief or in this anger through a gesture of friendship, they remain the grief and anger of my friend Paul ... And finally, if we undertake a shared project, this shared project is not a single project, and it is not presented to me and to Paul from the same angle; we are not equally committed to it, or at least not committed to it in the same way, from the mere fact that Paul is Paul, and I am myself. As much as our consciousnesses construct through our own situations a common situation in which they communicate, it is nevertheless from the background of his own subjectivity that each projects this “single” world.’ (Merleau-Ponty 2013, 414)

If the aim, in defending an account of plural subjecthood, is to defend the possibility of ontologically collective phenomena, then this passage could not be a clearer statement of how Merleau-Ponty does not view his account of the interaction between bodily subjects as opening up to this possibility. Even if we take part in a shared project, this can never be a literal sense of sharing for Merleau-Ponty, in which the project could be understood as that of a single subject. Although we share in a certain cultural and bodily world, in which we can meet on equal terms and as subjects, my

subjecthood and your subjecthood remain mutually exclusive forces that do not submit to a genuine intermingling.

To put it differently, as seen earlier, Merleau-Ponty, in trying to come to terms with this encounter, is concerned with how the 'I' might be understandable in the mode of the 'you' (i.e. the other 'I') or in the mode of the 'one' (i.e. the generalised or anonymous 'I'), but he is not concerned with how it might be grasped in the mode of the 'we' (i.e. the pluralised 'I'). For Merleau-Ponty, others do not present a problem for me and I am intimately familiar with them through our shared cultural world, as well as the shared bodily spaces that we inhabit, but they are nevertheless *other* and my familiarity with them presupposes this otherness. As such, it is within the categories of self and other that the theory must remain, for they are its very ground.

Conversely, what is being pursued here is an account of plural subjecthood, which requires an explanation of how this very subjecthood might be shared. As argued in Chapter 2, our experiences of thinking and acting together ought to be characterised as plural-yet-unified. This is to say, if 'we' is ever used in a non-distributive sense, it cannot be to refer to another supra-individualistic 'I'. In other words, the *we* must lie somewhere between a set of loosely associated individuals and another, spatially disparate, supra-individual in which the subjecthood of its constitute individuals has dissolved entirely. The very idea of a plural subject is that of a subject that is simultaneously a construction of the individuals that constitute it and more than the sum of its individualistic parts. As such, while an account of the plural subject must start from the presumption of individual subjecthood (without whom there is nothing to pluralise), it must also be able to explain how these individual subjectivities might be transcended. As an account of intersubjectivity, not only is this not what Merleau-Ponty is trying to present, but it is actively what he is trying to avoid. He requires an account of how two subjects might meet one another on equal terms and yet remain distinct individual subjects.

Undoubtedly, Merleau-Ponty's accounts of the extended habitual body and intersubjectivity move us on from a conception of the subject as a discrete or self-enclosed entity, wholly distinct from the world in which it finds itself. For Merleau-Ponty, there is a sense in which the subject almost dissipates into the world as a sort of unfinished project towards that world. Nevertheless, this very possibility is structured by the fact that each of us is, first and foremost, a distinct individual subject. Ultimately, therefore, the first-person perspective remains sacrosanct for him and is the underlying assumption of the particular way in which the world and others disclose themselves.

IV: Plural Objective Self-Awareness

While Merleau-Ponty's own account cannot provide us with a view of how subjecthood might be pluralised, predominantly because this was not his concern, it will be suggested that what has been said so far on the nature of the bodily subject does provide us with all of the necessary tools to develop just such an account ourselves. In other words, what Merleau-Ponty's understanding of the bodily subject and its interaction with others does contain is the raw materials out of which a theory of plural forms of subjecthood can be developed, even if he was himself inclined to see the subject as an essentially discrete entity.

In order to see how, let us return to the distinction that has structured a great deal of the discussion thus far. It has been argued, extensively, that one has two broad forms of self-awareness available: awareness of oneself as subject and awareness of oneself as object. While Merleau-Ponty's conception of the bodily subject breaks down the division between the two (the significance of which will be seen shortly), and renders the subject on a seamless spectrum from the most subjective to the most objective form of self-awareness, this is not to deny that there is any distinction to be made. Moreover, an underlying assumption of the way in which this distinction has been explicated above is that, at least on the Zahavian conception of the subject, there is a sense in which subjective self-awareness is ontologically prior to objective self-awareness. This is firstly because subjective self-awareness is constitutive of subjecthood in its most basic or minimal form; to be a subject, first and foremost, is to be aware of oneself as subject, which is to be given to oneself as the pre-reflective subject of experience. Secondly, and partly in virtue of this, it is because on this account of the subject it would be incomprehensible that there could be any reflective awareness of self unless there was already given, pre-reflectively, a subject on which to reflect. In other words, to have oneself as object presumes that one is already given to oneself as subject.

It is this priority of subjective self-awareness, it seems, that motivated Schmid to turn to the idea of experiences with a plural subjective character. For, intuitively, if we want to defend the possibility of plural subjects, and we believe that subjective self-awareness is the minimum requirement for there to be a subject at all, then we are led to posit the idea that groups must be capable of possessing some form of subjective self-awareness in their own right. Such a plural subjective self-awareness being thought of as constitutive of the most basic or minimal form of plural subjecthood. In other words, Schmid attempts to mirror Zahavi's approach to individual subjecthood in the plural case.

However, it should now be acknowledged that this approach, that of mirroring the individual case, involves a certain kind of error. For a start, as was made clear above, there is a sense in which the

plural subject is dependent on the individual subject, as there can be no plural subjects unless there are already individual subjects available for pluralisation. As such, while the individual subject can be thought of as a sort of primordial entity, not requiring a prior form of subjecthood or self-awareness in order to come into existence, the plural subject must be thought of as directly dependent on the individual. More importantly, however, in light of the conception of subjective self-awareness as naturally individuating through its intrinsic bodiliness, it must now be recognised that if plural subjective self-awareness is simply meant to refer to subjective self-awareness had by a group of individuals as opposed to a single individual, then the idea is, strictly speaking, incomprehensible.

This will seem immediately confusing, as it has been stated repeatedly that the aim of this thesis is to defend a conception of plural subjective self-awareness. However, the point is not that this idea itself makes no sense, but that it cannot be thought of as on a par with individual subjective self-awareness. For, if subjective self-awareness is naturally individuating through its intrinsic bodiliness, then this is necessarily the case. In other words, to be subjectively self-aware just *is* to be an individual subject; it is for there to be experience that presents itself, at the barest level, as from a particular spatio-temporal location, and it is hard to understand how such a location could possibly be understood as harbouring some sort of inbuilt plurality. Thus, the for-me-here-ness of experience must be thought of as synonymous with unity, singularity, and individuality. There is no room within it to accommodate both the thought that it can be constitutive of a single individual bodily subject and the thought that it might be constitutive of a subject constructed out of a plurality of such individual subjects. It is precisely such subjective self-awareness, understood in terms of for-me-here-ness, which is denied to the group. Rather, all subjective self-awareness is to be thought of as individual, and for a group of individuals to possess such a form of self-awareness would be for them to cease to be distinct individuals and instead become a single supra-individualistic entity.

However, while groups cannot be thought of as subjectively self-aware in the same sense as individuals (that is, unless they are some sort of Borg-like entity), groups can be thought of as objectively self-aware in spite of their lack of subjective self-awareness. While it is correct to hold that subjective self-awareness must come before objective self-awareness in the individual case, this need not be the order of things in the plural case. In fact, it is precisely because the individual subject must pre-date the plural subject that allows for the possibility of plural objective self-awareness. For, what distinguishes the individual from the plural subject is that in the former case there must first be some basic form of self-awareness in order for there to be a subject in the first place, but in

the latter case there is already self-awareness and subjecthood from which to build, that is, the self-awareness and subjecthood of the constitutive group members. Self-awareness does not need to be put on the table, as it were, because it is already there. As such, the group as a whole can be thought of as objectively self-aware in virtue of each of the individual members of the group being aware of themselves as constituting a particular group.

In order to understand what is meant by this claim, let us consider the example of a group of football players brought together to form a team in order to compete in a tournament. Perhaps the group numbers some fifteen people in total, including substitute players, and let us suppose that they do not know each other particularly well; all causal friends and acquaintances, but who have never actually played football together before. Each member of the team, having been told that they have been selected for the squad and knowing the other players who have also been selected, can quite accurately think of the group in question in terms of its particular members. This is to say, that they can each take the group as intentional object defined around the task in question, namely, 'those who are to try and win this tournament'. In doing so, it can be said that each of the members of the team are objectively self-aware of the team as a whole, understood in an ontologically distributive sense. For, each has taken as their object a group of which they form an intrinsic part: 'those who are to try and win this tournament', where each is one of 'those'. To put it differently, each of the players is part and parcel of the thing of which they are aware, and if they were not part of the team in question then their respective awarenesses would be of a very different sort, that is, they would be experiences of 'them' as opposed to 'us'. Thus, it would be misguided to think that somehow this form of awareness is not a kind of *self*-awareness.

Moreover, it would be equally misguided to think that the awareness in question must reduce down to an awareness that the individual players have of themselves as performing a certain role in the group, in conjunction with some awareness of the respective roles of all the other players. This objection would have it that, while the football team does display a form of objective self-awareness, the self of which each of the team members of aware is simply their own individual subjecthood, but held in some sort of relation to the other team members. Of course, such a situation is possible. For, each of the players in the team can, and probably does quite frequently, think of their own individual role in the team without having to pay much thought to the team as a whole. After all, it would be impossible for a football team to compete if the respective players had no sense of their specific roles. A striker must know that their primary task is to score goals just as a defender must know that their role it is to prevent goals being scored. Nevertheless, it would be absurd to think that it is not also possible for each of the individuals to think of the group as a whole, understood as

an object constituted by the particular team members; that it is somehow impossible for the players to have a sense of the team. Once this point is accepted, there appears to be no reason to think that the reflective awareness that each of the individual team members has of the group is not a form of objective self-awareness.

The sceptical reader might object that the supposed objective self-awareness in question seems like a different *kind* of awareness to that which is more commonly enjoyed by individuals, in the respect that in the group case that of which the individual is reflectively aware involves other selves. And, indeed, there is some substance to this claim. In fact, the basis on which the individual can be objectively self-aware is a question that has been mostly glossed over in the course of this thesis. What has been argued is that objective self-awareness is part of the very same spectrum of experience as subjective self-awareness. But, while this might demystify the relation somewhat, it does not provide a complete answer to the question, 'why are some of the objects of which I am aware, me?' However, putting this question to one side and assuming that, on whatever basis, we do enjoy forms of objective self-awareness, it must be recognised that some of the objects of which I am objectively self-aware will not relate to myself qua individual, but to myself qua group member. Of course, the important difference between such objects and those in the individual case, is that part of the properties of a group, understood as an object, are the individual subjects who constitute it. But there is no reason to think that this fact alone should rule out the individual members of the group forming a reflective awareness of the group and this awareness being a form of self-awareness. After all, on Merleau-Ponty's view of things, we have no problem forming a conceptual understanding of other subjects, so there appears to be no reason why we cannot be self-aware of an object that intrinsically includes other subjects. The only significant point is whether or not the individual team members' understanding of the group is presented to them as first-personal.

Having accepted that the group can form a part of the objective self-awareness of each of the individual team members, it is a small step to the claim that there can be a form of plural objective self-awareness. For, all it is for there to be plural objective self-awareness is for the individual members of a group to take such a reflective stance towards the group as a whole. This is to say, that I can form an idea, and thereby reflect upon, the group of which I am a part, understood as a distributive collection of individual subjects. And, provided that all of the other members of the group also form such an awareness, or at least would be compelled to if they stopped to consider the group (which of course, just as in the individual case, they might not), then it can be said that the group as a whole is objectively self-aware. For, all of those who constitute the group, namely,

the individual members, are aware of the group and plural objective self-awareness is to be understood as nothing more than the summation of the reflections of its members.

Unlike in the individual case, where there must already be a subject in some sense before that subject can become aware of itself as a particular object in the world, in the plural case a pre-existing plural subject is not necessary for a group to form an awareness of itself as a particular object, for the individual members of the group are able to form this awareness themselves. To put it differently, in the course of the explication of the bodily subject (in Chapter 5), intentional awareness was likened to a kind of light. Wherein, to be aware of oneself as object is simply to direct this light back on oneself and attend to the contours of one's being (contours that are inevitably transformed in virtue of being attended to). In the case of groups being aware of themselves as objects, this light can similarly be thought of as being directed back toward the self. But in this case the light is that of all of the members of the group, directed back upon that which they together constitute (i.e. the group as a whole).

If this still strikes the reader as controversial, it is crucial to note that to suggest that the group can be aware of itself in this way, is in no way to suggest that it can constitute a plural subject in any non-distributive sense. In many ways, what is under discussion is a particular form of intersubjectivity; to assert the possibility of plural objective self-awareness is to posit a specific relation that can hold between distinct individual subjects and the awareness that they can form of the groups that they inhabit. The members of the group are only aware of the group as the summation of the individual subjects that constitute it and in virtue of that which holds them in common. To put it differently, to say that a group can be objectively self-aware in this way is to say that the individual members can be aware of themselves as a particular determinate or concrete group. In much the same way as being objectively self-aware, in the individual case, is to be aware of oneself in terms of certain properties, for a group to be objectively self-aware is for the individual members to be aware of the group in terms of its properties, where the primary properties of any group are the particular individual subjects that constitute it and make it the particular group that it is, as opposed to some other group with a different membership. Thus, one does not have to go beyond a merely distributive or individualistic understanding of groups in order to accept that they are capable of being aware of themselves as a particular kind of object. For a group to be objectively self-aware is simply for the members of that group to be aware that there is a group of which they are a member.

V: The Group and Habituation

With this in mind, two features of the self-awareness available to groups can now be asserted. On the one hand, it is straightforwardly the case that groups can possess a form of objective self-awareness simply in virtue of the members' ability to take the group as a whole as intentional object. However, as emphasised, for them to take the group as object in this way is not for the group to progress beyond a merely distributive form of self-awareness. On the other hand, a group cannot, by definition, possess subjective self-awareness, at least in the sense in which it applies to individual subjects. To be subjectively self-aware in this sense *is* to be an individual subject. What is required of an account that wishes to defend the possibility of plural subjects, therefore, is an explanation of how the self-awareness of groups might occur between these two extremes.

It is at this point that Merleau-Ponty's account of the bodily subject as existing on a spectrum demonstrates its significance with regard to the plural subject. For, as was shown in Chapter 5, subjective and objective self-awareness should not be viewed as distinct species of experience, but rather two poles of the same spectrum. It is not that I can either take a reflective (objective) or pre-reflective (subjective) stance on myself. Rather, my awareness of self comes in infinitely various shades from the more reflective to the more pre-reflective in virtue of the structure of experience itself. As such, suggesting that groups can be objectively self-aware does at least place the group on the same spectrum of being as more unified forms of subjecthood. The problem is simply to understand how this self-awareness might progress beyond the most objective extreme of this spectrum.

Moreover, as Merleau-Ponty also demonstrates, our body schema, as a subjecthood lived in a cultural world inhabited by others, is enacted in shared bodily spaces. Thus, to be a bodily subject is, in part, to share the spatiality of one's body with others; to come into contact with one another by means of our bodily subjecthood as this subjecthood carves its way through the world. For Merleau-Ponty, of course, this sharing should not be understood as a sharing of subjecthood. Nevertheless, by starting from the premise that groups can possess a form of objective self-awareness, we can view this primordial intersubjective bodily sharing as the means by which individuals participating in groups can be involved in processes of habituation that would allow them to move beyond the objective pole of self-awareness, toward more unified forms.

On the account discussed above, habituation is that process whereby the subject incorporates the world into the spatiality of its own body schema. It is that by which the subject constitutes themselves as more than a simple or abstract being and enlarges the sphere of their power toward

the world. In other words, one is first a subject in a more abstract or minimal sense (in the sense of for-me-here-ness) and then stretches one's being out into the world by incorporating that world into one's body schema through habit. However, the members of a group already possess such schemas; they are already stretched out into a world that they intimately share with others and in which they can be objectively self-aware of the groups of which they are members. Allowing this, and allowing that objective self-awareness is one end of the same spectrum on which subjective self-awareness lies, we can think of processes of habituation at the group level, not as the means by which the relevant individuals can extend the purview of their own body schemas, but as the means by which those individuals can begin to form a familiarity with one another such that they start to develop a single shared body schema. In other words, group level habituation can be thought of as the means by which the group as a whole can move from an objective, and therefore distributive, form of self-awareness, towards increasingly subjective forms of self-awareness, which would constitute them as a single unified subject. It can be thought of as the process by which the group moves *from plurality toward unity*.

In order to grasp what is meant by such a unifying process of habituation, let us return to the example of the football team. As suggested above, they might begin their efforts to win the tournament together with only an objective, and therefore distributive, sense of the group that they constitute. They all know who the team consists of and the single task in which they are all thereby involved, but they have no prior familiarity of playing football with one another. However, as they begin to train and play together, this will change. At first they may each start to familiarise themselves with the playing patterns of the other players. Perhaps the central midfielder realises that if she picks up possession in an attacking position then the winger on the right tends to play on the overlap and that she can, therefore, often pick the winger out with a through ball. Perhaps the striker becomes aware that if the right winger picks up the ball down the flank, that most of the time she will swing in a cross and she ought, therefore, to gamble with a run into the box. Over time, and by becoming increasingly familiar with the way in which they each like to play (probably involving a lot of communication at first and a certain amount of complaining when things don't go as they would like), these forms of awareness will begin to recede into the background. The central midfielder might, at first, have to think actively about the playing tendencies of the right winger in various scenarios and, as such, pay a great deal of attention to the winger's movement. But the more that the two players begin to synchronise with one another's behaviour, the less that this will need to be the case. After some time, the central midfielder will probably have what we might call an instinctive, but more accurately a pre-reflective, awareness of what the right winger will do and

vice versa. She won't have to focus or pay attention to either what the right winger is doing or is likely to do. Rather, this awareness will structure how she sees the situation as a whole and allow her to focus on aspects of the game that she cannot take for granted in the same way, such as the less familiar and predictable behaviour of their opponents. In other words, the players will become habituated to one another's playing styles and this will allow them to function more effectively as a successful football team.

However, unlike in the individual case, these processes of habituation should not be thought of as leading to an extension of the individual body schemas of each of the players. It is not that the central midfielder brings aspects of the right winger into the spatiality of her own body. After all, as we have seen above, we do not confront one another as objects in space. Rather, we encounter others immediately in virtue of their bodily subjecthood and the fact that this subjecthood is, therefore, already out in the open. Thus, the players do not confront one another as tools or instruments to be subsumed within their bodily being, but as subjects with their own unique body schemata and spatiality⁶⁷. As such, in becoming habituated to one another's playing styles, this must be understood not as the process by which they extend their own individual body schemata, but by which they progress toward a single shared body schema. In many ways, it is precisely the sovereignty of the individual, that which was described above as the sticking point in Merleau-Ponty's position, that makes sense of such a potential sharing of bodiliness. For, it is in virtue of the fact that the individual members of the team recognise one another's subjecthood, that they cannot be understood as subsuming the other within their own body schemas, for the other is not an object to be subsumed.

Consequently, habituation in such a scenario must lead to a sense of what *we together* both do and constitute. In other words, when the central midfielder and the right winger's awareness of one another's playing habits recedes towards the pre-reflective level, what they will be left with is an experience of the game, not in terms of how they play as individuals in conjunction with their understanding of how the other plays the game, but in terms of how they, as a somewhat unified entity, play the game. To put it differently, their individual behaviour will be structured by their increasingly pre-reflective, and therefore unified, sense of how the team thinks and acts as a team; a sense of *how we play football*.

⁶⁷ This is not to say that it is impossible for them to treat one another as objects, because intuition and experience would suggest that there are those who do regularly treat one another as such. Nevertheless, all that is required for the present argument is that the reader accepts that we do often, and probably more often than not, recognise the subjecthood of others through their bodiliness and that, as such, treating others as an object to be subsumed within our own bodily spatiality is, if possible, the exception and not the rule.

What Merleau-Ponty shows us is that experience, and therefore the subject constituted by such experience, exists in between subjective and objective poles. It is in the milieu between the two in which subjecthood is lived. What Zahavi's conception of the subject shows, in light of this, is that the individual extends out from the subjective pole to the objective. Of course, as Merleau-Ponty emphasises, their subjecthood cannot be distinguished from the world in which it is constituted (and indeed it is this fact that renders the subject intrinsically bodily). But as a subject, we first know ourselves as an almost pure subject of embedded experience, before we understand ourselves as a concrete and discoverable thing in the world in which that experience is embedded. All that has been added to this is the idea that the plural case works in the opposite direction. The group knows itself firstly as a concrete and discoverable thing in the world. It knows itself through the awareness that its members have of that which they together constitute, and it is upon this that they build a growing sense of unified subjecthood through habituation. So, while the individual knows itself first as a minimal subject and then extends into the world through habituation, the group knows itself first as object and then habitually retracts, or perhaps withdraws, towards a unified minimal subjecthood

Of course, the contours of this subjecthood, that is, the shared body schema of, for example, the football team, will often be vague, ambiguous, and open to fluctuation, but this ought to be viewed as a strength of the position rather than a weakness. For a start, we have already seen that such features are also true of individual bodily subjecthood as its spatiality is affected by its own processes of habituation. But more importantly, it is precisely this fact that allows this position to account for the plural-yet-unified character of our experiences of thinking and acting together. For, the members of the team do not form a supra-individual or Borg-like entity. While a pre-reflective sense of what they do and how they carry themselves, as a football team, toward the task at hand is central to their success, it is equally central, as was noted above, that the individual players are able to think individualistically a great deal of the time. Moreover, as is clear from the above explication, while there will be some overarching sense of self that applies to the team as a whole, there will also be subsets of players who form particularly strong habitual bonds, and the history of football is laced with examples of such playing partners who have formed what might seem from the outside to be an almost telepathic connection with one another. In this vein, for a team to progress towards a more subjective form of self-awareness will, inevitably, involve a complex hierarchy of such plural subjective self-awarenesses enjoyed by subsets of players. However, the crucial point is that while a group can progress toward the extreme pole of subjective self-awareness, this is still a form of awareness that is, in its purest form, denied to them. A team can become increasingly subjectively

self-aware, and in doing so progress from a distributive to an ontologically collective sense of self, but what they cannot achieve is a truly singular and unified perspective on the world; the experiences of the group can never be characterised by a pure and simple for-me-here-ness. As argued above, for their experience to present itself in this way would be for them to cease to be a plurality of subjects altogether and become, instead, an individual. It is, therefore, the group's movement from their being a distributed plurality of individuals toward their being the single unified subject of experience that captures the paradox of their being at once plural and unified, for they must be understood as always somewhere between the two; always on the road to supra-individuality, but unable to ever reach this destination.

This is not to say that there won't be forms of plural subjecthood that will get closer to the supra-individual end of the spectrum. While the football team is an apt example in the respect that it emphasises processes of shared bodily habituation, it is unlikely that a football team will ever achieve a level of unification that progresses much beyond the distributive, if only for the reason that the membership of the average football team changes so regularly. In this respect, long term relationships between two people would seem to be the most obvious place to look for forms of plural subjecthood that approach the supra-individual. Not only is the group in such a context a stable set, but there is also a sort of inbuilt motivation towards harmonisation. Part of what it is to form and maintain such a bond with another human being is to grow together, and in doing so come to form a shared set of values and perspectives. Much like *inosculation*, in which the branches or trunks of two trees fuse or graft to one another as they grow in close proximity, so those in a long term relationship move from merely having similar points of view to literally sharing a single point of view on various matters. For instance, the well-worn cliché of 'being able to finish one another's sentences' would seem particularly significant in this context. Anyone who has had the privilege of the company of an elderly couple who have spent the majority of their lives together, will be familiar with the remarkable extent to which two people can come to rely on one another and form, what looks to all intents and purposes to be, a single interlocutor in a conversation. It is also worth noting that the feeling of expecting one's partner to finish one's thought in such a scenario only for them to express a completely different view on the matter, is actually curiously close in character to the disconcertion, discussed earlier, that would be accompanied in reaching for one's ear and not finding it. Nevertheless, regardless of how close such couples might come to achieving subjective self-awareness, it would be absurd to think that they have ceased to be individuals altogether and have instead become a Borg-like entity. For sure, there is less plurality and more unity on show than the average football team, but plurality remains nonetheless.

VI: Objections and Responses

It is undoubtedly the case that there will be certain immediate objections to this formulation of the idea that groups can achieve a certain level of subjective self-awareness (from here on I will take *plural subjective self-awareness* to refer to whatever this level might be), so it is worth trying to pre-empt a few of the more obvious challenges that might be made.

One such objection might be that there can be nothing akin to the immunity to error through misidentification in the case of groups, despite the fact that it was claimed in Chapter 3 that such immunity of self-reference is to be explained in terms of the non-ascriptive nature of subjective self-awareness. This is, therefore, suggestive of the fact that groups cannot be subjectively self-aware. To remind the reader, such immunity relates to the fact that there are certain statements made in the first person which do not appear to be open to a certain form of error. For instance, if I make the claim ‘I believe that x’, this statement might be mistaken, but it cannot be mistaken because I have attributed the belief to the wrong subject; it cannot be that I misidentified myself as the subject of that belief, when in fact it belongs to someone else.

It seems undeniable that no such immune first-personal statement will be available in the case of groups. This is to say, that no-one is ever in a position to state ‘we believe that x’, where it is inconceivable that the statement might be mistaken in virtue of the fact that the speaker misidentified the subject of the belief. It will always be a possibility in such cases that the speaker thinks that they are speaking on behalf of the group, when in fact they are mistaken. Perhaps they might even be mistaken that any such group exists. A good example of this was already discussed above, in which one of the partners in a long term relationship expresses a view that takes the other off-guard; they thought that they were on the same page, as it were, but as it turns out they were not. As long as such cases remain a possibility, then it appears as if statements made in the first person plural will always be open to error through misidentification. In fact, Schmid (2014c) has recently explored this fact in some detail, under the rubric of first person authority, and he too is led to conclude that such authority does not appear to be available in the group case. The only scenario that might, he thinks, provide such a possibility is that in which a group of individuals uses a communicative device such that they are able to simultaneously express the statement in question. Even if possible, however, he suggests that this is going to be an incredibly rare and peculiar circumstance.

Nevertheless, there is a very simple response that can be given to this objection, which is to point out that there is simply no reason to think that such immunity should be exhibited in the group case,

at least on the basis of the account of plural subjective self-awareness advanced above. For, as suggested, immunity to error through misidentification is a feature of self-reference that arises on the basis of the *non-ascriptive* nature of subjective self-awareness, in its most basic form. As has been argued, however, it is precisely this pure or extreme form of subjective self-awareness that is denied to groups; the possession of such self-awareness should be thought of as synonymous with individual subjecthood. To put it differently, there is no reason to think that statements made in the first person plural should be immune to error through misidentification, because the group is not a subject in the same wholly unified sense as the individual. There is no single subject that can speak with unanimous authority for the group, because the group is not a single subject in any simple sense; it is as subjectively complex as the processes of habituation that have brought it into being. In a sense, therefore, there cannot be immunity to error through misidentification in the group case simply because there is no-one who could even potentially make such a judgement. An individual cannot state with certainty what the attitude of the group is, because the individual is not the group. Equally, the group cannot state with certainty what the attitude of the group is, because it lacks the unity necessary from which to either make or express such a judgement.

So much for the idea of immunity to error through misidentification. A number of more substantial worries might arise once we begin to consider how the idea of plural subjective self-awareness can be applied to various test cases. The question being whether or not it is able to accommodate all of those circumstances in which it would seem most apt to ascribe plural subjective self-awareness to a group.

One such case is offered by, what we might call, more social animals. Let us take the example of a pack of wolves. On the one hand, intuition would suggest that if any group is going to possess a form of plural subjective (i.e. pre-reflective) self-awareness, then it is going to be such a pack, whose collective hunting habits are crucial for their survival. On the other hand, and this is where the problem lies, intuition would also suggest that the kind of awareness that is going to be denied to the pack is a form of plural objective self-awareness. This is to say, it is hard to imagine that the individual wolves are able to adopt a reflective stance towards the pack as a whole, at least if this is understood in anything like the same sense as in the human case. It is a fairly safe assumption that the capacity to consider oneself in terms of a complex range of properties is going to require a cognitive capacity that is unique to humans. However, it is claimed above that plural subjective self-awareness involves a group first having a form of plural objective self-awareness and then undergoing processes of group habituation. If this is the necessary order of things, then it would

appear to rule out the possibility that the pack of wolves, and groups of animals in general, have plural subjective self-awareness⁶⁸.

One option would be to bite the bullet and accept that plural subjective self-awareness is a purely human phenomenon. However, this seems particularly unpalatable when part of the appeal of Zahavi's account of the subject is that it is based on a form of self-awareness that will be possessed by any creature that has phenomenal experience. If, for instance, we were to understand thinking or acting together as resting on the ability to form peculiar we-mode intentions, as Searle and Tuomela do, then this would seem like a potential basis on which to distinguish between the capacities of humans and non-human animals. Indeed, as discussed in Chapter 2, something like this thought seems to be at the heart of Tomasello's work; that such intentions represent a developmental stage that only humans can achieve. But if we are to understand the ability to think and act together as resting on a plural subjecthood that is constituted by the pre-reflective character of experience, then there seems to be no reason to think that certain animals won't obtain such a plural subjecthood. It is entirely plausible that the wolves' behaviour in the midst of the hunt is going to be largely shaped by a pre-reflective awareness of how the pack hunts.

As such, we are forced to explain how the plural subjective self-awareness of a pack of wolves might come about in a way that is not exactly the same as in the human case. One way in which this might be achieved is by arguing that plural subjective self-awareness does not *always* need to develop out of an objective form of self-awareness. A plausible suggestion is that, perhaps, certain animals are hardcoded to think in such ways; that instinct simply drives the wolves to think from the perspective of the pack. However, by ruling out a prior form of plural objective self-awareness, such an approach would have the consequence of suggesting that the pack can be subjectively self-aware in a way that humans cannot. This is to say, that without a process of development from a more objective form of self-awareness, it would imply that the pack can experience from a single unified spatio-temporal perspective; that a pack of wolves is a Borg-like supra-individualistic entity.

A more convincing approach is to appeal to a middle-ground between such hardcoding and the kind of group habituation that occurs in the human case. This would rely on the thought that, while the wolves cannot achieve the kind of objective self-awareness enjoyed by humans, self-awareness is a spectrum from the most subjective to the most objective, and the self-awareness of the individual wolves can at least progress somewhere along this spectrum, even if it cannot extend very far beyond the most subjective pole. In other words, a wolf can at least have a very minimally objective

⁶⁸ My thanks to Tom McClelland for pointing this problem out to me.

sense of self. Once such a minimal objective self-awareness is ascribed to the wolves, then the kind of hardcoding discussed above, and which it is not unreasonable to think might play some role, can be thought of as allowing the kinds of leap in plural self-awareness that might not be possible, or at least will be more difficult, in the human case. This is to say, that a mixture of habituation and hardcoding may allow the wolves to jump from a minimal plural objective self-awareness to a quite robust and unified plural subjective self-awareness. This is a highly speculative suggestion, but it offers at least a reasonable view of how the experiences of groups of non-human animals might be captured by the same framework of plural subjective self-awareness, without portraying them as supra-individualistic entities.

It is also the case that, even when it comes to groups of humans, the process of habituation, as it was spelt out with regards to the football team, should be viewed as a sort of idealisation. The football team example is useful because it provides what is a relatively simple scenario in which the relevant stages of habituation are emphasised. But there is no reason to think that, even in the human case, the process will always be so linear.

Recognising this provides a response to another potential objection, in the form of the claim that it is unclear how the above account will deal with the fact that many of our experiences of thinking and acting together, even those which seem to exhibit the greatest degrees of unity, are often fleeting. For instance, it is quite often the case that a footballer will join a team and fit seamlessly into the existing group dynamic, being able to very quickly form the kind of bonds with other players that are necessary for success. And such *instant bonds* or *connections* seem to be a fairly regular feature of life in general, to a greater or lesser extent; it often seems that we are able to think or act in concert with others without such thought and action having to be the product of extensive familiarisation. And it is also the case that these instant bonds are often broken as quickly as they were formed. If correct, plural subjective self-awareness would seem to have to be a phenomenon that is at least capable of rising and falling with a swift regularity. An account that relies on complex, and perhaps quite lengthy, processes of habituation looks as if it is going to struggle to account for this fact.

However, the point is that the case of the football team that develops a sense of unity over time, through complex processes of habituation, is merely illustrative of a more general process that is going to be altered, complicated, and simplified by a number of factors. For a start, it seems plausible to assume that something like the hardcoding present in the case of the pack of wolves is also going to be present in the case of humans, if not quite to the same extent. This might be through certain genetic predispositions towards sociality, but more likely it will be the product of the

extensive socialisation that is undergone as an infant. Moreover, and partly as a result, in many cases there will be a certain interchangeability of the subject in cases of plural subjective self-awareness, just as there is an interchangeability of the object in the case of individual bodily habituation. To recall the case of the blind man and his cane, it is not that he becomes habituated to a particular cane, and thereby brings it within the unique spatiality of his body, but that he becomes habituated to the use of canes that are similar to the one with which he is familiar. This is to say, it is not the particular determinate object that is significant. And, perhaps more importantly in the present context, his habituation to one type of cane will almost certainly make it much easier for him to become habituated to the use of other types of cane as well. It is in this way that we should think of the new member of the football team who manages to seamlessly mesh with the existing group dynamic. The existing football team is not habituated to the use of particular determinate football players, as it were, but to football players *such as these*. If a football player joins the team who plays in a way that just happens to resemble the style of the player that she replaces, it should come as little surprise that the team will not have to undergo the same lengthy process of habituation all over again. There is, in other words, room within this account to accommodate the idea of fulfilling a certain role or place in the group. And, while this is only one respect in which processes of group habituation might be affected, it is at least indicative of the malleability that can be accommodated within the account.

Nevertheless, a final potential counter-example, for which there doesn't appear to be such an easy response, is offered by the case of *massive groups*. Consider, for instance, the partisan crowd watching our football team, which, owing to the rapid success of the team, now numbers in the thousands. In such a case it seems that the experiences of the crowd will have a shared affective or emotional component; provided that the team is winning, they are likely to enjoy a strong feeling of shared elation. Of course, one might want to argue that all that is shared in such a case is the type of emotion, rather than a certain subjective character of that emotion, but that does not seem quite capable of capturing the phenomenon in question. After all, the experience of watching the football team would be radically different if one was the only person in the crowd, and part of the appeal of watching football is the almost tribal experience of losing oneself in the group (a fact that also explains the inherent danger of such crowds). If this is accurate, then just like the example of the wolves, we would appear to be confronted with a group that, intuitively, looks like it will have a

pre-reflective form of self-awareness (i.e. an affective engagement with the match that is presented as *for-us*), but which doesn't involve any process of habituation⁶⁹.

Unlike the case of the wolves, however, here the response should be open yet sceptical. In principle I don't want to deny the possibility that there might be space within the theory presented above to account for such forms of experience. If there are plural subjects then, unlike individual subjects which have a certain uniformity, they will be as varied in form as the variety of shapes and sizes that groups commonly assume. As Schmid says of plural self-awareness in general:

'[I]t may be very wide and thin, such as the "sense of 'us'" among members of mankind, or narrow and deep, such as in the specialized, but close relationship between partners in a business, who have been working closely together on a daily basis for years but would never go out for a beer together.' (2014b, 22)

However, I am also conscious of not trying to use the idea of plural subjective self-awareness to explain too much, as it were. It has been repeatedly emphasised above that there is an important distinction to be made between intersubjectivity and plural subjecthood. While the emotional experience of being part of a football crowd can, evidently, be of a very powerful nature, there is no reason to assume that this experience cannot be captured by appeal to intersubjective concepts such as emotional contagion. There are many forms of social experience that are of a particularly radical or extreme nature, but there is no reason to assume that they will all be a form of plural subjective self-awareness. The latter, after all, might often be of a quite subtle or discrete nature. Accordingly, the reader should not be under the impression that the success of the theory in question rests on its ability to capture all social phenomena and, while Schmid might be correct that there is some sense of *us* that is common to all humanity, one need not assume that the sense in question is constitutive of humanity as a plural subject in any sense. The concern is, instead, only with a subset of social phenomena, in which the boundaries between individual subjects themselves begin to breakdown, and I am sceptical that this is what is necessarily occurring in the case of massive groups.

Finally, a completely different tactic might be to object that there is nothing contained in the kinds of examples discussed above that rules out an ontologically distributive interpretation of what is going on. Of course, an ontologically collective interpretation has been offered by suggesting that they provide examples of cases in which groups of individuals can be seen to develop, or be in the process of developing towards, a form of plural subjective self-awareness, but there is no intrinsic feature of such examples that prevents them from being interpreted otherwise by the individualist.

⁶⁹ My thanks to Manos Tsakiris for pointing this problem out to me.

Indeed, there is a sense in which this is a fair point. While some of the thoughts above do seem to tell against such an individualistic interpretation (for instance, such an interpretation would appear to involve the assertion that individual football players are *only* capable of encountering one another as tools or instruments to be brought under the purview of their own body schemata), there is nothing that rules out such an interpretation on anything like *a priori* grounds. However, to think that this is or should be the purpose of the argument explicated above is itself misguided. As was suggested from the offset, the aim of this thesis has been to argue against individualism by placing ontologically collective forms of thought and action on as sound a theoretical foundation as individual forms. In this respect the account above is successful just in case it presents a view of ontologically collective phenomena that is coherent and reasonable. For, as long as it achieves this then it has demonstrated that, at the very least, the idea of such phenomena cannot be ruled out as a possibility. An argument in favour of this view of plural subjecthood that did want to present a conclusive defence of the idea that we do, in actuality, think and act in this manner, would have to include a serious consideration of empirical evidence. After all, it is clearly the case that we very often do think and act individualistically even when operating as a member of a group, and empirical evidence would be required to show that the situation is ever more collective than this. But, similarly, such evidence would also be required on the part of the individualist in order to demonstrate that the situation is *never* more collective than this. As it stands, both are viable ways in which we might think and act as part of a group, and this alone is already enough to show that at least some of the assumptions of the individualist are mistaken.

What is of the upmost importance, and the key criteria on which the above account ought to be judged, is that the conclusion that has ultimately been reached should not grate against the reader's intuitions regarding the nature of their experiences of thinking and acting as part of a group. For, it was suggested in Chapter 1 that the intention of this thesis was to take an idea, that of groups attaining a form of phenomenal self-awareness, which undoubtedly strikes one's intuitions as radical on first glance, and dissolve this radicalism. In this respect, while the arguments that have led us to this point might have forced the reader to re-evaluate some of their ideas regarding subjecthood and the nature of experience, there should be nothing that is particularly unintuitive above the conclusion itself, in part because it recognises the important respects in which plural subjecthood and plural subjective self-awareness will differ from individual subjecthood and individual subjective self-awareness.

VII: Final Conclusions

At the beginning of this thesis it was suggested that there appears to be a guiding assumption throughout virtually all of the literature on contemporary forms of anti-individualism. This assumption is that the idea that groups could be phenomenally self-aware to be so absurd that any defence of ontologically collective forms of thought or action will have to avoid precisely such a view of things. If groups have minds, or are subjects, or can act, then this will have to be because such features are either not dependent on phenomenal self-awareness or are fundamentally different in the group case from the individual. However, having now explored some of these anti-individualist arguments in more detail, it could be suggested that this assumption actually encompasses two distinct intuitions: (1) If groups are phenomenally self-aware then the ‘what it is like’ of this group self-awareness will have to be identical to the ‘what it is like’ of individual self-awareness; (2) We have such a clear conception of the ‘what it is like’ of individual self-awareness, that we can rule out in advance the possibility of self-awareness being extended to groups.

In one way or another, the argument of this thesis has now demonstrated that both of these intuitions are, in fact, mistaken. With respect to (2), as phenomenological thought in general makes crystal clear, describing the nature of awareness and capturing the way in which the subject appears to itself, is an extraordinarily complicated task. Moreover, once we begin to unravel the *Gordian Knot* that is the subject’s encounter with the world, we discover a range of ideas that are actually quite conducive to theorising in ontologically collective terms. This, in essence, is the *additional argument* of this thesis described in Chapter 1. By beginning with Zahavi’s thought that the subject is first constituted by a form of subjective self-awareness, this position has been gradually added to, such that what we are left with is a conception of the subject that is rich in a variety of ways. On this account to be a subject, even in the most basic or minimal sense, is to already be individuated by the bodiliness of subjective self-awareness. Moreover, it is to be intrinsically tied to more concrete forms of being through the spectrum of bodily self-awareness, and formed through processes of habituation such that the world is open to incorporation within the spatiality of one’s body. Thus, most importantly, to be a subject is not to be hermetically sealed from the world, but to be naturally extended into, and open to, a world that is shared with other bodily subjects. It is only once we understand the subject in this way that we uncover the possibility that there might be forms of plural, as well as individual, self-awareness.

However, and with respect to (1), it has also been argued that there is no reason to think that this means that we must understand the processes that lead to the formation of the individual subject, so constituted, as simply carrying over to the plural case. On the contrary, it has been shown that

we can more accurately grasp the nature of plural subjecthood, and the plural subjective self-awareness on which it must rest, as being formed by processes that flow in the opposite direction. After all, what must be true of any form of plural subjecthood is that it presupposes the subjecthood of the individuals that constitute it. Thus, whereas in the individual case the subject ought to be understood as first subjectively self-aware and then developing itself as a more or less concrete being in the world via processes of habituation, in the plural case the group must be understood as first objectively self-aware in a purely distributive sense and then progressing towards a unified subjective sense of self via processes of group habituation.

It is on this basis alone that ontologically collective forms of thought and action become a viable possibility for groups. They may never be able to achieve the same kind of unified phenomenal self-awareness that is enjoyed by individuals, but they can nevertheless attain more than a merely distributive sense of self and, in doing so, they can begin to form the kind of unified subject which can both think and act. In Chapter 1 it was suggested that the statement ‘Tim and Tom are lifting the piano’, assuming that it is meant in the linguistically collective sense, admits of two interpretations; one in which Tim and Tom form the single unified subject of the action and one in which the statement is actually a way of referring to two distinct actions of two distinct subjects (albeit ones that add-up to the lifting of the piano). However, it should perhaps now be said that there are, in fact, a range of more ambiguous interpretations available between these two alternatives and it is precisely this ambiguous terrain that is the realm of the plural subject; the subject that is at once plural *and* unified.

It is in virtue of this ambiguity that the account of plural subjecthood on offer does not provide neat answers to many of the questions that academics have sought to answer within the literature on contemporary anti-individualism. For instance, one of the central questions in that discussion has been whether or not groups can be held morally responsible. Such an interest is understandable considering how often we do apportion moral culpability to groups (the reaction to BP in the wake of the oil spill in the Gulf of Mexico, is a very clear demonstration of this fact). However, if this line of questioning rests on the assumption that individuals can be held morally responsible in virtue of the fact that they are singular unified agents, it should be clear that the above account of plural subjecthood does nothing to diffuse the controversy of assigning moral responsibility in the group case. It is precisely the extent to which groups can be unified that has been shown to be rich in complexity. Not only can a plural subject never be unified in the same sense as an individual subject, and not only will it never be a simple matter to determine the extent to which a group has moved toward a unified subjective sense of self, but even then there are likely to be complex hierarchies

of plural self-awareness in play such as in the example of the football team. In other words, if assigning moral responsibility to someone or something is a matter of determining the extent to which they are a responsible unified agent, then this will never be a straightforward matter in the case of plural subjects.

Nevertheless, that this account presents such a complex response to such questions should be thought of as a strength rather than a weakness, because it is surely at odds with our intuitions regarding the nature of groups to think that a simple response either could, or should, be given. Groups are, more often than not, complex entities and, as such, any account that wants to suggest that groups can be held morally responsible must be capable of capturing this complexity. Thus, the account on offer is successful in the respect that it provides a complex answer to a complex question. Groups may well be capable of being held morally responsible, but determining if and when they can will be as complex as the variety of shapes and sizes in which plural subjects appear.

Finally, in acknowledging this complexity we ought also to recognise a question that has only been dealt with briefly above: the question of the types of group that this account might be applied to. In the previous section some remarks were made as to the kind of groups that might, and might not, be explicable in terms of plural subjective self-awareness, but very few definite conclusions were reached. The reason why this question has not been dealt with in more detail, despite the fact that in many ways it is of central importance to any such account, is because there is a sense in which the position presented above is purposely abstract. As suggested above, the point has not been to present a theory that makes sense of all group formations, but it has been to present a theory that has the flexibility to be applicable to a broad range of situations. As should be clear from what has been stated above, there are groups that will find it easier to achieve more unified forms of subjective self-awareness than others. For instance, a couple in a long term relationship will enjoy a simplicity and stability that will make it relatively easy to undergo quite robust forms of group habituation and build a single shared body schema over an extended period of time. On the other end of the scale, large and complex groups, such as companies, will likely exhibit a high degree of hierarchical complexity as well as a lack of long term stability in membership, which will make it far more difficult to achieve such self-awareness. However, it is precisely the purpose of the account of ontologically collective thought and action on offer that it ought to make sense of these differences. Just as bodily self-awareness comes in a spectrum from the most subjective to the most objective, so the awareness enjoyed by groups ranges from the most objective, and therefore distributive, to the most subjective, and therefore unified.

Bibliography

- Anscombe, Elizabeth. 1983. "The Causation of Action." In *Knowledge and Mind: Philosophical Essays*, edited by Sydney Shoemaker and Carl Ginet, 174–90. Oxford University Press.
- Bacharach, Michael. 1999. "Interactive Team Reasoning: A Contribution to the Theory of Co-Operation." *Research in Economics* 53 (2): 117–47.
- . 2006. *Beyond Individual Choice: Teams and Frames in Game Theory*. Edited by Natalie Gold and Robert Sugden. Princeton, N.J.: Princeton University Press.
- Bermúdez, José Luis. 2013. "Bodily Awareness and Self-Consciousness." In *The Oxford Handbook of the Self*, edited by Shaun Gallagher, Reprint edition, 157–179. Oxford; New York, NY: Oxford University Press.
- Bratman, Michael E. 1987. *Intentions, Plans, and Practical Reason*. Cambridge: Harvard University Press.
- . 1999. *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge: Cambridge University Press.
- . 2009. "Shared Agency." In *Philosophy of the Social Science: Philosophical Theory and Scientific Practice*, edited by Chrysostomos Mantzavinos, 41–59. Cambridge: Cambridge University Press.
- . 2014. *Shared Agency: A Planning Theory of Acting Together*. New York, NY: Oxford University Press.
- Brook, Andrew. 2001. "Kant, Self-Awareness and Self-Reference." In *Self-Reference and Self-Awareness*, edited by Andrew Brook and Richard C DeVidi, 9–30. Amsterdam; Philadelphia: J. Benjamins.
- Carman, Taylor. 2008. *Merleau-Ponty*. London; New York: Routledge.
- Carr, David. 2012. "Cogitamus Ergo Sumus: The Intentionality of the First-Person Plural." In *Interpreting Husserl: Critical and Comparative Studies*, 281–96. Springer.
- Cassam, Quassim. 1998. "Introspection and Bodily Self-Ascription." In *The Body and the Self*, edited by Jose Luis Bermudez, 311–36. Cambridge, Mass.: MIT Press.
- . 2013. "The Embodied Self." In *The Oxford Handbook of the Self*, edited by Shaun Gallagher, Reprint edition, 139–56. Oxford; New York, NY: Oxford University Press.
- Clark, Andy, and David Chalmers. 1998. "The Extended Mind." *Analysis* 58 (1): 7–19.
- Elster, Jon. 1985. *Making Sense of Marx*. Cambridge; New York: Paris: Cambridge University Press.
- Evans, Gareth. 2001. "Self-Identification." In *Self-Reference and Self-Awareness*, edited by Andrew Brook and Richard C DeVidi, 95–141. Amsterdam; Philadelphia: J. Benjamins.
- Ewing, Ac. 1996. *A Short Commentary on Kant's Critique of Pure Reason*. Chicago: University of Chicago Press.
- Fitzpatrick, Dan. 2003. "Searle and Collective Intentionality: The Self-Defeating Nature of Internalism with Respect to Social Facts." *American Journal of Economics and Sociology* 62 (1): 45–66.
- Frankfurt, Harry G. 1988. *The Importance of What We Care about: Philosophical Essays*. Cambridge University Press.
- Gallagher, Shaun. 1998. "Body Schema and Intentionality." In *The Body and the Self*, edited by Jose Luis Bermudez, 225–44. Cambridge, Mass.: MIT Press.
- Gallagher, Shaun, and Dan Zahavi. 2007. *The Phenomenological Mind: An Introduction to Philosophy of Mind and Cognitive Science*. London; New York: Routledge.
- . 2015. "Phenomenological Approaches to Self-Consciousness." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2015.
- Gallotti, Mattia, and Chris D. Frith. 2013. "Social Cognition in the We-Mode." *Trends in Cognitive Sciences* 17 (4): 160–65.

- Geach, Peter Thomas. 1957. *Mental Acts Their Content and Their Objects*. Routledge and Kegan Paul.
- Gilbert, Margaret. 1989. *On Social Facts*. Princeton, N.J.: Princeton University Press.
- . 1990. "Walking Together: A Paradigmatic Social Phenomenon." In *The Philosophy of the Human Sciences*. Vol. XV. Midwest Studies in Philosophy. University of Notre Dame Press.
- . 1993. "Agreements, Coercion, and Obligation." *Ethics* 103 (4): 679–706.
- . 1998. "In Search of Sociality." *Philosophical Explorations* 1 (3): 233–41.
- . 2008. *A Theory of Political Obligation: Membership, Commitment, and the Bonds of Society*. Oxford: Oxford University Press.
- Gold, Natalie, and Robert, Sugden. 2006. "Introduction." In *Beyond Individual Choice: Teams and Frames in Game Theory*, by Michael Bacharach, 1–34. Princeton, N.J.: Princeton University Press.
- . 2007a. "Theories of Team Agency." In *Rationality and Commitment*, edited by Fabienne Peter and Hans Bernhard Schmid, 280–312. Oxford: Oxford University Press.
- . 2007b. "Collective Intentions and Team Agency." *The Journal of Philosophy* 104 (3): 109–37.
- Henry, Aaron, and Evan Thompson. 2013. "Witnessing from Here: Self-Awareness from a Bodily versus Embodied Perspective." In *The Oxford Handbook of the Self*, edited by Shaun Gallagher, Reprint edition, 228–49. Oxford; New York, NY: Oxford University Press.
- Hume, David. 1888. *A Treatise of Human Nature*. Edited by Lewis Amherst Selby-Bigge. Oxford: Clarendon Press.
- Hutchins, Edwin. 1996. *Cognition in the Wild*. Cambridge, Mass.: MIT Press.
- Hutto, Daniel D., and Erik Myin. 2013. *Radicalizing Enactivism: Basic Minds without Content*. Cambridge, Mass: MIT Press.
- Kant, Immanuel. 1999. *Critique of Pure Reason*. Edited by Paul Guyer and Allen W. Wood. Cambridge; New York: Cambridge University Press.
- Kriegel, Uriah. 2003. "Consciousness as Intransitive Self-Consciousness: Two Views and an Argument." *Canadian Journal of Philosophy* 33 (1): 103–32.
- . 2009. *Subjective Consciousness: A Self-Representational Theory*. Oxford: Oxford University Press.
- Krueger, Joel. 2013. "Merleau-Ponty on Shared Emotions and the Joint Ownership Thesis." *Continental Philosophy Review* 46 (4): 509–31.
- Legrand, Dorothee. 2006. "The Bodily Self: The Sensori-Motor Roots of Pre-Reflective Self-Consciousness." *Phenomenology and the Cognitive Sciences* 5 (1): 89–118.
- . 2007a. "Pre-Reflective Self-as-Subject from Experiential and Empirical Perspectives." *Consciousness and Cognition*, Subjectivity and the Body, 16 (3): 583–99.
- . 2007b. "Subjectivity and the Body: Introducing Basic Forms of Self-Consciousness." *Consciousness and Cognition*, Subjectivity and the Body, 16 (3): 577–82.
- . 2010. "Subjective and Physical Dimensions of Bodily Self-Consciousness, and Their Dis-Integration in Anorexia Nervosa." *Neuropsychologia*, The Sense of Body, 48 (3): 726–37.
- . 2013. "Phenomenological Dimensions of Bodily Self-Consciousness." In *The Oxford Handbook of the Self*, edited by Shaun Gallagher, Reprint edition, 204–27. Oxford; New York, NY: Oxford University Press.
- Legrand, Dorothee, and Susanne Ravn. 2009. "Perceiving Subjectivity in Bodily Movement: The Case of Dancers." *Phenomenology and the Cognitive Sciences* 8 (3): 389–408.
- Lukes, Steven. 1978. *Individualism*. Oxford; New York: Blackwell.
- Meijers, Anthonie W. M. 2003. "Can Collective Intentionality Be Individualized?" *American Journal of Economics and Sociology* 62 (1): 167–83.
- Merleau-Ponty, Maurice. 1964. *Sense and Non-Sense*. Reprint edition. Northwestern University Press.
- . 1969. *The Visible and the Invisible*. Evanston: Northwestern University Press.

- . 2008. *The World of Perception*. London; New York: Routledge.
- . 2010. *Child Psychology and Pedagogy: The Sorbonne Lectures, 1949-1952*. Translated by Talia Welsh. Evanston, Ill: Northwestern University Press.
- . 2013. *Phenomenology of Perception*. Translated by Donald A. Landes. Routledge.
- Noë, Alva. 2006. *Action in Perception*. Cambridge, Mass.: MIT Press.
- O'Shaughnessy, Brian. 1998. "Proprioception and the Body Image." In *The Body and the Self*, edited by Jose Luis Bermudez, 175–203. Cambridge, Mass.: MIT Press.
- Pacherie, Elisabeth, and Jérôme Dokic. 2006. "From Mirror Neurons to Joint Actions." *Cognitive Systems Research*, 7 (2–3): 101–12.
- Persson, Ingmar. 1999. "Awareness of One's Body as Subject and Object." *Philosophical Explorations* 2 (1): 70–76.
- Petersson, Björn. 2007. "Collectivity and Circularity." *The Journal of Philosophy* 104 (3): 138–56.
- Roth, Abraham. 2013. "Shared Agency." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2015.
- Rovane, Carol. 1998. *The Bounds of Agency: An Essay in Revisionary Metaphysics*. Princeton, N.J: Princeton University Press.
- . 2004. "What Is an Agent?" *Synthese* 140 (1-2): 181–98.
- Schmid, Hans Bernhard. 2009. *Plural Action: Essays in Philosophy and Social Science*. Dordrecht; New York: Springer.
- . 2014a. "The Feeling of Being a Group: Corporate Emotions and Collective Consciousness." In *Collective Emotions*, edited by Christian von Scheve and Mikko Salmela, 3–22. Oxford: Oxford University Press.
- . 2014b. "Plural Self-Awareness." *Phenomenology and the Cognitive Sciences* 13 (1): 7–24.
- . 2014c. "Expressing Group Attitudes: On First Person Plural Authority." *Erkenntnis* 79 (9): 1685–1701.
- Schmitt, Frederick F. 2003. "Socializing Metaphysics: An Introduction." In *Socializing Metaphysics: The Nature of Social Reality*, edited by Frederick F. Schmitt, 1–38. Lanham, Md: Rowman & Littlefield.
- Schweikard, David P., and Hans Bernhard Schmid. 2013. "Collective Intentionality." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2013.
- Searle, John R. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.
- . 1990. "Collective Intentions and Action." In *Intentions in Communication*, edited by Martha E. Pollack, Jerry L. Morgan, and Philip R. Cohen, 401–16. Cambridge, Mass: MIT Press.
- . 1995. *The Construction of Social Reality*. London: Allen Lane.
- . 1997. "Responses to Critics of The Construction of Social Reality." *Philosophy and Phenomenological Research* 57 (2): 449–58.
- . 2006. "Social Ontology: Some Basic Principles." *Anthropological Theory* 6 (1): 12–29.
- . 2007. "Social Ontology and the Philosophy of Society." In *Creations of the Mind Theories of Artifacts and Their Representation*, edited by Eric Margolis and Stephen Laurence, 3–17. Oxford; New York: Oxford University Press.
- . 2010. *Making the Social World: The Structure of Human Civilization*. Oxford: Oxford University Press.
- Shoemaker, Sydney. 1968. "Self-Reference and Self-Awareness." *The Journal of Philosophy* 65 (19): 555–67.
- . 1988. "On Knowing One's Own Mind." *Philosophical Perspectives* 2 (January): 183–209.
- . 1994a. "Self-Knowledge and 'Inner Sense': Lecture III: The Phenomenal Character of Experience." *Philosophy and Phenomenological Research* 54 (2): 291–314.
- . 1994b. "Self-Knowledge and 'Inner Sense': Lecture II: The Broad Perceptual Model." *Philosophy and Phenomenological Research* 54 (2): 271–90.

- . 1994c. "Self-Knowledge and 'Inner Sense': Lecture I: The Object Perception Model." *Philosophy and Phenomenological Research* 54 (2): 249–69.
- Simmel, Georg. 2011. *Georg Simmel on Individuality and Social Forms*. University of Chicago Press.
- Smith, Joel. 2006. "Bodily Awareness, Imagination and the Self." *European Journal of Philosophy* 14 (1): 49–68.
- Strawson, P. F. 1964a. *Individuals: An Essay in Descriptive Metaphysics*. London: Routledge.
- Thalberg, Irving. 1984. "Do Our Intentions Cause Our Intentional Actions?" *American Philosophical Quarterly* 21 (3): 249–60.
- Thompson, Evan. 2010. *Mind in Life*. Cambridge, Mass.: Harvard University Press.
- Tollefsen, Deborah Perron. 2006. "From Extended Mind to Collective Mind." *Cognitive Systems Research, Cognition, Joint Action and Collective Intentionality*, 7 (2–3): 140–50.
- Tomasello, Michael. 2009. "Why We Cooperate." MIT Press.
- Tomasello, Michael, Malinda Carpenter, Josep Call, Tanya Behne, and Henrike Moll. 2005. "Understanding and Sharing Intentions: The Origins of Cultural Cognition." *Behavioral and Brain Sciences* 28 (05): 675–91.
- Tomasello, Michael, and Hannes Rakoczy. 2003. "What Makes Human Cognition Unique? From Individual to Shared to Collective Intentionality." *Mind & Language* 18 (2): 121–47.
- Tuomela, Raimo. 1989. "Actions by Collectives." *Philosophical Perspectives* 3 (January): 471–96.
- . 1991. "We Will Do It: An Analysis of Group-Intentions." *Philosophy and Phenomenological Research* 51 (2): 249–77.
- . 1992. "Group Beliefs." *Synthese* 91 (3): 285–318.
- . 1995. *The Importance of Us: A Philosophical Study of Basic Social Notions*. Stanford: Stanford University Press.
- . 2000a. "Collective and Joint Intention." *Mind and Society: A Journal of Cognitive Studies in Economics and Social Sciences* 1 (2): 39–69.
- . 2000b. *Cooperation: A Philosophical Study*. Dordrecht; Boston: Kluwer Academic.
- . 2003. "The We-Mode and the I-Mode." In *Socializing Metaphysics: The Nature of Social Reality*, edited by Frederick F. Schmitt, 93–128. Lanham, Md: Rowman & Littlefield.
- . 2005. "We-Intentions Revisited." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 125 (3): 327–69.
- . 2013. *Social Ontology: Collective Intentionality and Group Agents*. New York, NY: Oxford University Press.
- Tuomela, Raimo, and Kaarlo Miller. 1988. "We-Intentions." *Philosophical Studies* 53 (3): 367–89.
- Velleman, J. David. 1997. "How To Share An Intention." *Philosophy and Phenomenological Research* 57 (1): 29–50.
- Zahavi, Dan. 2002. "First-Person Thoughts and Embodied Self-Awareness: Some Reflections on the Relation between Recent Analytical Philosophy and Phenomenology." *Phenomenology and the Cognitive Sciences* 1 (1): 7–26.
- . 2007. "First-Personal Self-Reference and the Self-as-Subject." *Consciousness and Cognition, Subjectivity and the Body*, 16 (3): 600–603.
- . 2008. *Subjectivity and Selfhood: Investigating the First-Person Perspective*. Cambridge, Mass.; London: MIT Press.
- . 2009. "Is the Self a Social Construct?" *Inquiry* 52 (6): 551–73.
- . 2013. "Unity of Consciousness and the Problem of the Self." In *The Oxford Handbook of the Self*, edited by Shaun Gallagher, Reprint edition, 316–35. Oxford; New York, NY: Oxford University Press.
- . 2015. "Self and Other: From Pure Ego to Co-Constituted We." *Continental Philosophy Review* 48 (2): 143–60.
- Zaibert, L. A. 2003. "Collective Intentions and Collective Intentionality." *American Journal of Economics & Sociology* 62 (1): 209–32.