

***Parts Across Space and Time* – Extended Book Proposal**

There are some central, foundational questions we must ask about the physical universe: questions about which parts objects have, and about how objects relate to the regions they occupy. These questions have implications throughout Metaphysics, including implications for which kinds of entities exist, how they persist across space and time, and how we should understand debates about these topics. This book is an introduction to and contribution to work on these topics. The central thesis is that, necessarily, any objects that extend across regions of time, space, or spacetime have proper parts contained within every disjoint region that is not free of the object. This thesis follows from a combination of arguments for (i) the impossibility of extended regions of space with no distinct regions making them up, (ii) the impossibility of *multilocation*, which involves an object being exactly located at more than one region, and (iii) the impossibility of *spanning*, which involves an object being located at (or filling) an extended, composite region without having any parts located at any proper subregions of the region. If extended, simple regions are impossible, then any object located at an extended region must extend across distinct regions. If objects cannot span extended, composite regions that they fill, it means that they must have parts located in each of the smaller subregions of the composite regions they extend through. If they cannot multilocate through the regions, then those parts located at those smaller subregions must be distinct from the whole object. So, if objects cannot span or multilocate, any object located at (or filling) an extended region must have proper parts within each composite region they fill. Because my arguments apply to temporal as well as spatial regions, they have direct implications for persistence through time.

Along the way to establishing my thesis, I cover much more. I present arguments about which theory of *parthood* we should accept, and what explanatory work we should expect our theory to do. I present arguments about how to correctly understand the top competing views of persistence, and argue in favor of one of them. And I present material that is at once critical of and helpful to philosophers who believe there is not a substantive debate about persistence. This book will serve as an introduction to Mereology and Location for those unfamiliar with the topics, providing a more thorough survey of the issues to accompany my recent edited volume, *Mereology and Location*. And the book will make new contributions to contemporary debates in Mereology, Location, Ontology, MetaOntology, and Persistence, proving valuable to researchers on these topics.

The table of contents is as follows:

Introduction

1. Theories of Parthood and Location
2. Extended Regions and Logics of Location
3. Problems for Multilocation and Mereology
4. Additional Problems for Multilocation
5. Problems for Spanners
6. How to be a Four-Dimensionalist
7. The Problem for Three-Dimensionalism
8. Implications for Deflationism

Introduction

I provide an overview of the book followed by a listing of assumptions and remarks on my methodology. Perhaps the most notable feature of my methodology is the extent to which I take unusual, sometimes impossible, cases seriously. There are two reasons for my doing this.

My first reason is simply that many of the views I consider are claims not just about metaphysical possibility, but about analytic possibility. Philosophers making claims about which systems of definitions to endorse (taking the definitions to correspond to real features of the world, capturing natural properties and relations) are making claims that they take to be analytically necessary. So, any analytically possible cases incompatible with their claims will be problematic for them, even if those cases are metaphysically impossible. For instance, Josh Parsons believes that what it means for an object to be located at a region is for the object to be present at every subregion within the region, and to not be present outside the region. Thus, if we can find an analytically possible case where an object meets those conditions for being located at a region but does not seem to be located there, we have reason to reject Parsons' account; it will seem that he has given us an incorrect definition, he hasn't captured *what it means* to be located somewhere.

My second reason for taking strange, impossible cases seriously is that I believe that necessary truths (with the exception of definitional truths or those about essences) tend to require explanation. So when we claim a case is impossible (that is: necessarily, it does not obtain), we must explain *why* it is impossible. This is in contrast with the view endorsed by many philosophers who think that the necessity of a claim often provides an explanation for that claim's truth, and that necessity itself does not require explanation. For instance, suppose that you flip a coin 100 times, and it comes up heads each time. This seems strange, and calls out for explanation. But now suppose you're told that not only does the coin come up heads every time it's *actually* flipped, in every possible situation where it could ever be flipped, it comes up heads then, too. That is, not only does it come up heads in the actual world, it comes up heads in every possible world. The philosophers who believe necessity is explanatory think that saying the coin couldn't possibly come up heads explains why it comes up heads every time you actually flip it; no further explanation is needed. In contrast, I think that the coin's coming up heads in *every possible world* is even more surprising than its happening to come up heads 100 times in a row when actually flipped. Far from needing no explanation whatsoever and being able to play an explanatory role, the fact that the coin comes up heads in every possible world is in

greater need of explanation (something about the coin being essentially weighted, for instance) than its contingent counterpart.

These higher expectations for explanation influence the sort of project we are engaged in. We are not just looking for accounts that are intensionally correct, doing “well enough” by happening to correctly describe features of a certain kind in all possible cases. I want to illuminate the explanatory relations between our claims. I illustrate this by looking at an example involving Weak Supplementation. Weak Supplementation is a central, widely endorsed claim in Mereology that says, roughly: if you have some but not all of something, there must be another part of it that shares no parts with what you already have. I have the reader consider two cases. In the first, we have something that looks exactly like a dinosaur, except that she has exactly one smaller part: a fleck of green scale. Weak Supplementation, correctly, tells us that this case is impossible: if the dinosaur has one part distinct from the whole, there must be some other part of the dinosaur that works with the first part to make the whole up. But we can construct a second case, where our dinosaur has exactly two small flecks of green scale, and no other parts distinct from it. We want to say this case is impossible, too: the two flecks of green scale are not enough to make up our dinosaur. But Weak Supplementation does not rule the second case out: each part distinct from the whole has another part to keep it company. What seems to be doing the work, then, is a Weak Supplementation of Pluralities principle: for any x s and y , if each of the x s is part of y , but the x s do not fuse to make y , there must be some part of y that has no parts in common with any of the x s. I make this argument in more detail in the text, discussing some complications involving Unrestricted Composition (UC). There are ways of combining UC with Weak Supplementation to get an intensionally correct Mereology. But I recommend adopting Weak Supplementation of Pluralities instead, because it is better at getting the *explanation* right with respect to these cases.

1. Theories of Parthood and of Location

This chapter provides the required background on theories of Parthood and Location. However, it includes new ways of understanding the material with the hope that this will provide further insight into the topics even for those thoroughly familiar with the issues. It overlaps somewhat with my introduction to the *Mereology and Location* volume I edited last year for Oxford University Press, but contains many explanations and arguments not present in that chapter.

I begin by working to make clear exactly what our aims should be in presenting theories of parthood and location. In giving any such theory, we must complete the following steps: (i) identify all of the relevant properties and relations, (ii) determine which properties and relations are defined in terms of which others, and (iii) list the formal rules governing the instantiation of the properties and relations (and how they can combine with relations of different sorts). There are further steps we can carry out once this is complete, such as (iv) listing the metaphysical constraints on the instantiation and recombination of these properties and relations, and (v) describing the explanatory work we take our theory to do.

I take the reader through most of these steps for each of Parthood and Location. And in doing so, I highlight a wealth of controversies that may otherwise be overlooked. I'll go into detail on just one of those below.

To complete the first step, we must identify all of the relevant mereological and locative properties and relations. Philosophers are typically content to do this by merely making a list. *Parthood*, *proper parthood*, *overlap*, etc., are included on the list, but *having a bunny-shaped proper part* is not, though that is about parthood as well. I claim that we can think of the relevant relations in relation to the following question: How much of x makes up how much of y ? We can give the answers on the following table (taking 'some' to mean 'some and not all').

How much of x makes how much of y ?

	Some of x	All of x
Some of y	Overlap	Proper parthood
All of y	Proper parthood	Improper parthood

The various ways of answering our question produce different mereological relations. We have overlap (part sharing), proper parthood (where x is part of y and distinct from y), proper parthood going in the opposite direction (where y is part of x and distinct from x), and improper parthood (the relation of an entity being part (and all) of itself). We'll notice that a crucial relation is missing, though: fusion, the relation of some things making up a further thing. Improper parthood might be said to be a special case of this. So we may amend our question:

How much of the x s make how much of y ?

	Some of the x s	All of the x s
Some of y	\approx Overlap	\approx Proper parthood
All of y	\approx Proper parthood	Fusion

Answers to this question give us our *fusion* relation, and it gives us *improper parthood* when there is exactly one of the x s. Similarly, we end up with notions roughly corresponding to *overlap* and *proper parthood*, and our original relations are special cases of these new relations, when there is exactly one of the x s. These new notions are more general than our original ones, and will apply to pluralities and groups as well as to individuals. In this spirit, we may amend our question once more, to produce even more general relations:

How much of the x s make how much of the y s?

	Some of the x s	All of the x s
Some of the y s	\approx Overlap	\approx Proper parthood
All of the y s	\approx Proper parthood	\approx Fusion

Thinking of parthood in this way, as just a special case of a more general, more basic relation, and taking that more general relation to be the subject of Mereology, gives us a

generality for Mereology that we did not have before. (This idea goes back to Stanislaw Lesniewski.) Mereologists often talk about how Mereology applies across ontological categories, and constraining and describing material as well as abstract objects, regions as well as occupants of those regions, and so on. This more general, more basic relation adds a layer to this, allowing us to have a single “making” relation that entities across all of these categories (individuals, pluralities, groups, etc.) can stand in, and we can describe the rules (transitivity, asymmetry, etc.) that the relation must follow in all cases. For instance, it is plausible that every kind of *making* relation is transitive: if the *xs* are some of what make up the *ys*, and the *ys* are some of what make up the *zs*, then the *xs* are some of what make up the *zs*. We also have room to debate about which rules, if any, the general relation must follow in exactly some cases.

I approach picking out locative relations in a similar way, as what we get with different ways of answering the question “How much of *x* is in how much of *r*?” I note that, just as we invoked a mereological “making” relation in asking our question about Mereology, we’re invoking a locative relation, “is in”, when asking our question about location. Which relations we produce will depend on which locative relation we invoke in the question we ask.

In the chapter, I cover the dominant views of which rules the relations must follow, and the reasons that have been given for challenging these rules. As I do so, I also present some of my own concerns about the definitions that have been given. My goal is not to argue for a new mereology or logic of location, but to make the reader aware of numerous points of controversy, and ways of challenging the most central, basic assumptions people widely make when using mereological and locative relations in their work.

2. Extended Regions and Logics of Location

This chapter borrows largely from my forthcoming “Placement Permissivism and Logics of Location”, in the *Journal of Philosophy*, as well as from my 2015 “Time-Travel and Fundamentality”, in *Thought*. In this chapter, I argue against the analytic possibility of extended, simple regions.

I note that there are three options for views about the mereological structure of regions: (i) we may think there are no simple regions (so either no regions exist, or regions are gunky), (ii) we may think that there are simple regions, and any simple regions are point-sized, or (iii) we may think that there are simple regions, and at least some of these are non-point-sized (and thus, on my use of the term, they are extended).

I argue that if we think extended, simple regions are possible, we can combine that possibility with other plausible claims to produce what I call “Place Cases”: cases where an object is present in an extended, simple region, but does not fill that region. The problem is that Place Cases are incompatible with every possible logic of location that has the following features: (a) it posits at most one primitive locative property or relation, and (b) it appeals to nothing beyond that property or relation, basic logic, and mereology in its definitions of all of the other locative properties and relations. I call such theories “parsimonious logics of location”, and every logic of location that has been endorsed in print is parsimonious. The issue is that parsimonious logics of location simply do not have

the resources to adequately capture all of the locative features of Place Cases: they do not have the resources to distinguish Place Cases and their less problematic counterpart cases with respect to locative features.

The Place Cases I discuss are all quite strange: the one I focus on, *Almond in the Void*, involves an almond contained in an extended, simple region that seems larger than the almond is. It is important to note, however, that even if we think this case is metaphysically impossible, this will not be enough to defend our parsimonious logics of location. Logics of location are, in part, about the *analytic* connections between various locative properties and relations. Thus, if we have even one analytically possible case where these connections fail to hold, we have a counterexample to our logic of location. It is not sufficient, therefore, to note that Place Cases are strange or even metaphysically impossible; we must claim that they contain a contradiction in terms, and this is much harder.

Here is how we might generate Place Cases: Much of the motivation for thinking that extended, simple regions are possible also supports the possibility of extended, simple regions of varying sizes. It is also plausible that if, possibly, there is a region of size s , there is some possible object that is at least that small. This gives us the possibility of an extended, simple region, and the possibility of an object that is smaller than that region.

We can apply recombination principles to these possibilities: for instance, consider David Lewis's *Patchwork Principle for Objects*. According to this principle, if there is some possible region of spacetime s , and some other (perhaps in a different world) possible region of spacetime, s' , there is some possible world where an intrinsic duplicate of s , containing duplicates of all of the inhabitants of s , is next to an intrinsic duplicate of s' , containing duplicates of all of the inhabitants of s' . Applied to our extended, simple region and our smaller object, this allows us to generate the possibility of a world where an intrinsic duplicate of the extended, simple region and an intrinsic duplicate of the smaller object are right next to one another. We can also apply a *Pushing Principle* of recombination to our objects. According to this principle, (a) if we have an object next to an empty chunk of spacetime that's at least as big as the object is, and there's (b) an empty path from where the object is to the region, and (c) the path is large enough for the object to continuously pass through, then it's possible to push the object from its starting location into the previously empty region, without making the object or the region undergo any intrinsic changes. This principle allows us to move from the possibility the *Patchwork Principle* allowed us to generate, to the possibility of an extended, simple region that contains an object that fails to fill it.

We also have the option of applying a more straightforward recombination principle, which I call *Possible Placement Permissivism*. It says: if an object bears some locative relation to some region, then for any possible, empty region of the same size and shape, an intrinsic duplicate of the object can bear the same location relation to an intrinsic duplicate of the empty region. So: if an almond in my refrigerator is contained within a one-foot cubed region, and possibly, there exists an empty extended simple region of the same size and shape, then possibly, an intrinsic duplicate of my almond is contained within an intrinsic duplicate of the extended simple region. This gives us the possibility of a Place Case.

Here's how Place Cases are incompatible with parsimonious logics of location. Consider *Almond in the Void*, in which there is an almond contained in an extended, simple region, the region is empty of everything except the almond and its parts, and the almond does not fill the region. Contrast this case with a less problematic counterpart case, which is a perfect match with respect to mereology, basic logic, and which regions the almond is weakly located in (i.e., which regions are not entirely free of the almond). This counterpart case will only differ with respect to the size of the extended, simple region, and anything directly following from that: in this case, the almond *does* fill the region, and the almond and the region are the same size. Call this case *Almond in its Shadow*. Intuitively, the Shadow and Void differ with respect to whether the almond fills and is exactly located at the region it is in: in the Shadow case, it fills its shadow and is exactly located at that region. In the Void case, the almond does not fill the region that contains it, and is not exactly located anywhere. My general argument is this: if we endorse a parsimonious logic of location, we get at most one basic locative relation. If we take that to be a relation that is held fixed between the Shadow and Void cases, we'll have to say that all the non-basic relations are the same between the cases, too. If instead we take our basic relation to be one that is instantiated in just *one* of the cases, then for any non-basic relation, if the non-basic relation is defined in a way that requires instantiation of the basic relation, the non-basic relation won't be instantiated in the other case. If instead it's defined more liberally, then we won't always get the right results about *which* regions objects bear the non-basic relation to.

If we attempt to respond to the incompatibility of Place Cases and parsimonious logics of location by adopting additional locative primitive, we will be faced with either inexplicable necessary connections, or unacceptable descriptions of possible cases. Consider: it is overwhelmingly plausible that, necessarily, if an object fills a region, then the object is weakly located at that region. With a parsimonious logic of location, we can explain this via appeal to the definition of *weak location* or the definition of *filling a region* (one must be defined in terms of the other, or they must be defined in terms of something in common). For instance, if to fill a region is to be weakly located at every subregion of that region, our necessary claim is an immediate consequence. If, instead, we take *weak location* and *filling a region* to each be primitive, we lack an explanation of why the instantiation of the latter requires instantiation of the former. We may have a host of these inexplicable, necessary truths. Alternatively, and even worse, we may simply claim that these aren't necessary truths after all, and it is possible for something to fill a region without being even weakly present there.

My recommendation is to endorse a parsimonious logic of location, and instead respond to my arguments by rejecting the analytic possibility of extended, simple regions. (I also discuss reasons for rejecting the analytic possibility of gunk without thoroughly discontinuous parts.)

3. Problems for Multilocation and Mereology

This chapter borrows largely from my 2011 "Multilocation and Mereology", in *Philosophical Perspectives*, and presents some of the central arguments for this book. However, in that paper I was unable to draw out the implications of the argument, and I

was unable to offer a positive view in response to the one I was attacking. Further, there is a related argument about fundamentality that I was unable to give, and which I discuss in detail in the next chapter of the manuscript.

Some philosophers believe that multilocation is possible: they believe that it is possible for a single entity to bear the *located at* relation to multiple regions. For instance, many philosophers believe that universals do this. Other philosophers believe that this would occur in time-travel scenarios (where, for instance, you travel back to visit yourself at your high-school graduation, you are located at a region in the bleachers, and located at a region on the stage). In this chapter, I argue that multilocation produces conflicts with some of the most widely accepted principles of Minimal Mereology: the Transitivity of Proper Parthood, the Asymmetry of Proper Parthood, and Weak Supplementation. In fact, because these principles are often taken to be analytically true, positing even the mere conceivability of cases involving various kinds of multilocation will require the denial of the analyticity of the axioms of Minimal Mereology.

Here is the central case: imagine you have a statue of a dog, made of smaller statues. Call the dog-statue Clifford. One of the smaller statues that makes it up is a statue of a dog biscuit, Kibble, which makes up most of Clifford's torso. Kibble is a proper part of Clifford. Kibble also has statues among its proper parts, including Odie, a solid dog-statue that makes up most of Kibble. Clifford is put outside in the sun for 20 years. In that time, 3 things occur: (i) Clifford melts a bit, becoming a solid, continuous mass. Unfortunately this results in Clifford no longer having any statues as proper parts; none of them survive melting together. But Clifford survives, since most of the same matter stuck around. (ii) Next, all of the atoms that once made up Clifford are slowly replaced by new ones. But this change happens very gradually. (iii) Finally, Clifford shrinks quite a bit, and many of his atoms leave for good. Over the course of 20 years, Clifford comes to be about a third of his original size. But again, the change is quite slow, and Clifford is resilient. One day, God looks down and says, "Clifford, you've been a good boy. I'm going to let you go on an exciting trip. I'm going to let you travel back in time to just before you were created." Perhaps Clifford is helped to move discontinuously, or perhaps spacetime is simply looped, providing a path to the past a dog-statue can follow. But one way or another, Clifford ends up sitting on a shelf, visiting this earlier time and occupying exactly the region occupied by Odie. Then an artist comes along and sees him and says, "I know just what to do with you!" The artist uses Clifford to make up most of Kibble, just like Odie. In fact, Odie and Clifford don't share just these properties, they share all of their properties: they're one and the same dog-statue! So Clifford is located within Kibble, which is located within Clifford.

Here is the problem: if Clifford is identical to Odie, and Odie/Clifford (at the small region) is a proper part of Kibble which is a proper part of Odie/Clifford (at the large region), then if proper parthood is transitive, Odie/Clifford is a proper part of itself. That is, this case requires that we either give up the transitivity of proper parthood, or claim that proper parthood is not irreflexive: something can be a proper part of itself. Since *proper parthood* is commonly understood as parthood with distinctness, this is a problem. (If we respond to the case by relinquishing the irreflexivity of proper parthood, the case will also require that we give up the asymmetry of proper parthood; we can no longer claim that, for any x and y , if x is a proper part of y , then y is not a proper part of x .)

I discuss several responses to these problems, and focus on the option of relativising parthood to regions, either to one region or to two. That is: instead of saying simply that x is a part of y , we say (following Hud Hudson) that x is a part of y at r , or (as Cody Gilmore and I have independently developed) that x at $r1$ is part of y at $r2$. I argue against these responses: if we replace the axioms of Minimal Mereology with their region-relativised counterparts, we will not be able to capture the intuitions that supported the original axioms. I walk the reader through the numerous problems for the view on which parthood is relative to a single region, discuss the ways in which relativising parthood to *two* regions might be thought to help. Then I show that even twice-relativised parthood is inadequate: it is unable to capture Weak Supplementation (and its variants, such as Weak Supplementation of Pluralities), and it is unable to correctly describe cases that involve a combination of colocation and multilocation. (I also discuss and argue against recent alternative accounts given by Karen Bennett and Aaron Cotnoir.)

If we agree that multilocation causes problems for Mereology in cases like these, and we want to avoid those problems, there are two ways we might respond. Our first option is to attempt to restrict Mereology and multilocation to different domains. The second option is to reject the possibility of multilocation altogether.

4. Additional Problems for Multilocation

In this chapter, I attempt to provide additional support for the general rejection of the possibility of multilocation in response to the Clifford/Odie case by presenting similar, related problems. The first is a problem related to fundamentality, the second is related to new versions of the Problem of Temporary Intrinsic.

In raising a worry for fundamentality, I note that the same cases we can use to raise problems for the transitivity, irreflexivity, and asymmetry of *proper parthood* can be used to raise problems for the transitivity and asymmetry of *more fundamental than*. This is because philosophers often think that differences in fundamentality often co-occur with instances of the parthood relation. Frequently, philosophers believe that proper parts are more fundamental than the wholes they're parts of (though some philosophers believe this is reversed, and that wholes are more fundamental than their parts). But then, if we have a case where something seems to be a proper part of itself (if parthood is transitive), this will also be a case where something is more fundamental than itself (unless *more fundamental than* is not transitive). Thus, our theorist is looking at an inconsistent quartet. She must say either (a) the case is not possible, (b) *more fundamental than* does not track proper parthood in this way, (c) *more fundamental than* is not transitive, or (d) *more fundamental than* is not irreflexive and it is not asymmetric.

I believe the best response is (a), denying the possibility (even the analytic possibility) of the case. However, I consider various ways of endorsing (b), and claiming that the connection between proper parthood and *more fundamental than* is slightly more subtle than we may have thought. I consider, for instance, the view on which only simples (things without proper parts) are more fundamental than things with proper parts. I present two worries for this, one of which is that we may think there are no simples at all; we might think the world is gunky, where everything has an infinite chain of smaller parts. In

such a world, there would be no “bottom level”, but we may nonetheless think the world contains entities that differ with respect to how fundamental they are. (I present similar worries for theorists who think the only things more fundamental than others are the entities at the “top level”.) There are ways of endorsing (b) that attempt to isolate multilocation problems; they may say, roughly, that *more fundamental than* tracks proper parthood *unless* multilocation is involved. I claim that this is implausible as well, describing a world containing nothing but multilocated entities; arguably, this world still includes differences in fundamentality.

This problem involving fundamentality, it might be thought, is even more pressing than the problem for Mereology, because some of the options we had available for Mereology are much less plausible for fundamentality. For instance, though we may respond to the problem for Mereology by relativising *proper parthood* to regions, it is less plausible to relativise *more fundamental than* to regions. And though we may have hoped to respond to the problems for Mereology by simply giving up the relevant axioms about *proper parthood* (replacing the transitivity of *proper parthood* with the mere transitivity of *parthood*), we cannot plausibly deny the analogous axioms about *more fundamental than* (it is not plausible to claim that our inclination to predicate transitivity of *more fundamental than* simply a confused inclination to predicate transitivity of *at least as fundamental as*. Those are not relations we tend to conflate).

I argue that, in light of these considerations, there are only two plausible responses to my Clifford/Odie multilocation cases: either restrict multilocation to a domain outside the scope of the rules we intuitively take to govern the parthood relation (and the fundamentality relation), or deny that multilocation is analytically possible. Only the latter will allow us to claim that the axioms of Minimal Mereology (and similar claims about fundamentality) are analytically necessary and universal.

Further, additional pressure to reject the possibility of multilocation comes from reflection on new versions of the generalised Problem of Intrinsic. The Problem of Temporary Intrinsic is familiar: I am thinking about Metaphysics at noon, I won't be thinking about Metaphysics at 3am (when I will be asleep). But this means that I'm thinking about Metaphysics, and not thinking about Metaphysics. I'm differing from myself, and this involves a violation of Leibniz's Indiscernibility of Identicals. Of course, we only get the result that I'm differing from myself if we ignore the role time is playing. The puzzle, though, is in figuring out exactly how time rescues us. Some people take properties to be relations to times; it is no worse for me to be philosophising-at t1 and not-philosophising-at t2 than it is for me to be taller than the podium but not taller than the building. Other philosophers take me to have different temporary parts at each time. Those parts, they say, are the things with the different properties: my current temporal part is thinking about Metaphysics, and later, a distinct temporal part will fail to think about Metaphysics, so nothing differs from itself. There is a generalised version of this problem, originally presented by Kris McDaniel, which arises if an entity can be multilocated at a single time. If I visit myself at my high-school graduation, and I'm standing on the stage and not standing in the bleachers, it seems that at a single time I'm standing and not-standing. The response here tends to be to take spacetime regions to play the same role times played in the original Problem of Temporary Intrinsic. Though some responses to

the original problem, such as the positing of irreducible tense, will not produce plausible, analogous responses in the new Problem of Intrinsic involving multilocation.

However, new cases have been presented to raise problems for using any sort of region-relativisation to respond to new versions of the Problem of Intrinsic. John Hawthorne has presented a case where a ghost-like entity (though we could use bosons with opposite spin instead) seems to collocate with itself, differing from itself at a single time and region. Another case, presented by Joshua Spencer, involves an extended simple statue occupying an extended simple region. The statue, Spencer stipulates, is red in the middle, and grey elsewhere, but has no uniformly red or uniformly grey parts. The statue seems to differ in colour, but we cannot appeal to distinct parts, times, or regions to explain the difference. And a third case, one of my own, involves an object that, in an ordinary situation, would seem to multilocate, but which happens to do so within an extended simple region. If one combines the possibility of multilocation with the possibility of extended simple regions (of the right sort) and endorses the right accompanying recombination principles, they can generate a case like this. If the object differs from itself within the region, we will be unable to explain this difference via appeal to different regions.

One option, in responding to these cases, is to simply reject their possibility. And with the third, we will have to choose between rejecting the possibility of extended simple regions, rejecting our recombination principles, or rejecting the possibility of multilocation. These considerations may help to tip the balance against positing multilocation, given the other arguments that have been presented.

5. Problems for Spanners

In this chapter, I present reasons for believing that spanners cannot exist. A spanner is an entity that fills an extended, non-simple region, but which does not have any parts (proper or improper) located at any proper subregion of that region. In other words, these are extended simples that do not multilocate within themselves. For instance, imagine: You have a very special chess piece. This chess piece is solid marble, and has a volume of one cubic inch. It occupies an extended region (which, presumably, has smaller subregions making it up). However, this chess piece does not have any smaller parts. It does not have a top half, or a base, or a two-dimensional curved surface. It extends across a composite region, but does not have any smaller parts (proper or otherwise) that make it up.

I begin by surveying arguments for extended simples in the literature, and presenting worries for these arguments. For instance, I discuss an argument recently put forward by Kris McDaniel and Raul Saucedo, which goes roughly as follows: Suppose, for reductio, that necessarily, whenever an object is located at a region, the object and the region have the same parthood structure. (So, for instance, if an object is simple, the region it is at is simple as well.) Thus, there is a necessary correlation between the relata of the *located at* relation: necessarily, one has a given parthood structure iff the other also has that parthood structure. But, say McDaniel and Saucedo, we should follow Hume in endorsing a ban on any unexplained necessary connections between the inessential, intrinsic properties had by the ontologically independent relata of fundamental relations. If

mereological features are intrinsic to objects and regions, and either objects or regions have these mereological features contingently, and objects and regions don't share any parts, then we have a necessary connection that is problematically unexplained in exactly the way the neo-Humean principle is intended to rule out. We cannot explain this necessary connection via appeal to the definition of *located at*, because it is primitive. We cannot explain it via appeal to property inheritance, because the properties are both intrinsic, and the bearers of those properties do not share parts with one another. We cannot explain the necessary connection by claiming that neither object could be otherwise, because one of the objects has the property inessentially. So, there must not be a necessary connection of this sort. The mereological structures of objects and their locations can differ. And, the argument continues, one acceptable way for this to occur is for an object to have no proper parts while its location does have proper parts. My response to this argument involves a refinement of the Humean recombination principle, and a claim that even the refined principle will overgenerate possibilities.

In addition to surveying the arguments in the literature for and against the possibility of extended simples, I will present an argument of my own against the possibility of extended simples. I claim that if we think that a case involving colocation and extended simples is possible, we will be pushed to adopt a new mereological *overlap* primitive (or something that gives similar results). The rough idea is this: if we think colocation is possible, we can imagine a case where two spheres partially spatially overlap in virtue of having proper parts that are colocated. That is: in this first case, we have partial spatial overlap without mereological overlap, via colocation of parts. We can contrast this with a second case, in which there is partial spatial overlap that involves mereological overlap (part sharing) rather than colocation of proper parts. In this case, two spheres partially spatially overlap in virtue of having a proper part in common. If we think that extended simples are possible, we can imagine a third case that is a variant on the first. In this case, we have partial spatial overlap between two spheres, just as before, but these spheres are extended and simple, rather than composite. In this case, though the objects would partially spatially overlap, they would not have colocated proper parts because they do not have proper parts at all. Instead, they might be made partly of colocated portions of matter. Now imagine a fourth case, which is a variant on the second case. Just as we moved from the first case to the third by simply changing whether the spheres had proper parts, we will move from the second case to the fourth by making the same change. Here, we have two spheres that overlap spatially, but which cannot share parts because they are simple. Still, this case seems to differ from the third case: in this case, the spheres seem to be sharing *something*. Perhaps, for instance, the spheres are partly constituted by some of the same matter. If we accept the possibility of extended simples together with the possibility of colocation, we will be able to construct cases like these, and will need some way of accounting for the overlap between the two spheres in the fourth case. There is not *merely* spatial overlap, as there is in the third case, and there is not a standard sharing of parts, as there is in the second case. So we must expand our ontology or ideology in some way to account for the difference; our standard mereological *overlap* relation, which simply involves sharing of parts, cannot do the required work. I claim that in light of the other arguments concerning extended simples, the best response we can give to this case is to reject the possibility of extended simples.

6. How to be a Four-Dimensionalist

This chapter has two central aims. The first is to establish how we ought to understand four-dimensionalist persistence through time. I will argue for a formulation that, it turns out, entails (and, on the assumption that ordinary objects are temporally extended, is entailed by) the kind of persistence through time that I have argued for in the earlier chapters of this manuscript. The second aim of the chapter is to discuss how four-dimensionalists ought to think about the Clifford/Odie case in order to minimise counterintuitive the consequences of her view when applied to that case. I'll describe each of these components of the chapter in more detail.

There has been a significant amount of disagreement about how to understand the Three-Dimensionalism/Four-Dimensionalism debate. Some take it to be a debate about location, others take it to be a debate about Mereology, still others take it to be a debate about size or existence-in-time. I mention reasons for taking Three-Dimensionalism to be either a view on which objects occupy temporally unextended regions, have temporally unextended size, or exist-in-time (*a la* Kit Fine). I go into detail on how to formulate Three-Dimensionalism in the following chapter. We can thereby set Three-Dimensionalism apart from Four-Dimensionalism and the Spanning View (on which things span across time), as both latter views require that temporally persisting entities have size that includes extension in the temporal dimension. But to distinguish Four-Dimensionalism and the Spanning View from one another, we must invoke Mereology: Four-Dimensionalism requires that temporally extended entities have proper parts, the Spanning View does not – in fact, it requires that temporally extended entities do not have temporal parts located at proper temporal subregions of the regions they persist through. But to isolate the view we have in mind when we discuss Four-Dimensionalism, we must specify which sorts of proper parts the view requires extended objects to have.

On canonical formulations of the view, Four-Dimensionalism requires the existence of proper temporal parts, which are fusions of all of the object present at a time. The disagreement is then about how abundant these temporal parts must be.

I argue that these formulations of Four-Dimensionalism have this flaw: they either require needless commitments about the structure of time, or they require needless commitments about how liberally entities decompose into parts. One response to this worry is to give up on providing a single, neutral formulation of Four-Dimensionalism, and instead opting for a family of four-dimensionalist views each catered to the other views (such as views about the mereological structure of spacetime) that we endorse. I take this to be unsatisfying: there seems to be something in virtue of which all of the views in this family count as *four-dimensionalist*. That is, there seems to be some persistence-related natural feature (or group of features) that they all share, in virtue of which we put them in the same group, that of *four-dimensionalist* views. If we can capture what this feature (or group of features) is, it seems we will have captured what it is to persist four-dimensionally.

I walk the reader through a series of formulations of Four-Dimensionalism, showing that each is unduly restrictive to Metaphysics that we want to take to be

independent of the view. After demonstrating the problems with the current formulations, I present an alternative formulation that avoids these problems. According to this version of Four-Dimensionalism, temporally extended objects needn't have proper *temporal parts*, that is, fusions of all of the object present that is at any particular times. Instead, objects must have proper *temporary parts*, which are proper parts of the object contained within times the object persists through. Of course, requiring the existence of *some* proper temporary parts is not enough: my formulation also requires that, for any time the object persists through, it has proper temporary parts contained within that interval. (There are some additional requirements as well.) This version of Four-Dimensionalism not only provides us with a stronger, more flexible version of the view of persistence, it also helps us see more clearly what is central to Four-Dimensionalism, and what is really being debated when we discuss how objects persist through time.

The second project in this chapter is to discuss how a four-dimensionalist ought to describe the Clifford/Odie case. In setting up my problematic Clifford/Odie case in chapter 3, I note that there are reasons to consider positing multilocation regardless of our view of persistence: denying that multilocation occurs has counterintuitive consequences. But I then go on to argue that multilocation is impossible. If we want to deny the possibility of multilocation, then, we have some work to do in responding to these reasons in favour of it. Luckily, I argue, the four-dimensionalist has a way to explain why we were inclined to posit multilocation, while providing an alternative theory that does not involve any entity being located at more than one region.

The motivation for positing multilocation in self-visitation time-travel cases is this: we do not want to claim that when the time-traveller visits himself, he then has two heads, four hands, and is significantly spatially gappy. If we posit multilocation, we can claim that the entity has one head on the stage, and one head in the bleachers. It is nowhere true of the entity that it has two heads. If we reject the possibility of multilocation, however, it is not clear how we can give a similar result. If we think that temporally persisting entities derive their temporary properties from their temporal parts, and we think that a time-travelling entity has a single, non-multilocated temporal part at the time when the self-visitation occurs, then we will have the counterintuitive result that the time-traveller has twice the number of heads we would have thought.

The key in responding to this, I claim, is to follow Ted Sider in saying that in time-travel self-visitation cases, there is a single temporal part (or cluster of temporary parts) present, but also two *stages* of the entity are present. And entities derive their temporary properties from their *stages*. Thus, whereas we had a single thing of the right size twice-over in the case of multilocation, by appealing to stages we get *two* things of the right size to be the primary bearers of temporary properties

We can also distinguish between stages and temporal parts in fission cases. Consider: suppose Clive is cloven into a right half and a left half, each of which is paired up with a new half that is an intrinsic duplicate of the half it is replacing, producing 2 post-fission intrinsic duplicates of our pre-fission Clive. We can ask: does Clive survive? If we think he does, we may be undecided about which post-fission parts constitute temporal parts of Clive: is the temporal part multilocated? Is it spread across the regions filled by the 2 duplicates of Clive? However, in spite of not being decided on that issue, we can

easily pick out the person-stages (without taking a stand on *whose* person stages they are). There is a person-stage occupying the region occupied by one of the Clive duplicates, and there is another person-stage occupying the region filled by the other Clive duplicate.

The difficult project is in giving an account of what stages are. I begin by discussing and rejecting Sider's account, which involves appeal to sortals. I then look at initially tempting accounts that appeal to immanent causation, and accounts that appeal to co-parthood relations. I conclude by returning to appeals to sortals, and combining such an appeal with claims about grounding to produce an intensionally adequate account. I remain neutral on whether this captures the essential features of stages, but I offer this as one way a four-dimensionalist may avoid positing multilocation while retaining intuitive results in time-travel cases.

7. Three-Dimensionalism and the Rejection of Multilocation

In this chapter, I argue that insofar as my cases provide pressure for rejecting the possibility of multilocation, they also provide pressure for rejecting Three-Dimensionalism combined with Eternalism.

The case for this will be straightforward if Three-Dimensionalism in combination with Eternalism requires multilocation. And, without positing any new primitive properties or relations, it seems to. To present some support for this claim (though I recognise arguments from elimination are problematic), I look at candidate formulations of Three-Dimensionalism that do not involve appeal to location. For, if it is possible for a three-dimensionally persisting entity to be wholly present at a time without being located at a spatial region that is entirely present at that time, then there must be some way of giving an account of *wholly present* that does not involve appeal to location. The best candidates involve appeal to mereological features. As I examine these candidate formulations of Three-Dimensionalism, I present counterexamples to each. That is, for each formulation, I present at least one case that it misclassifies with respect to whether it involves three-dimensional persistence.

For instance, here is one candidate formulation of Three-Dimensionalism that I consider and reject:

- *Three-Dimensionalism_{Mereology}*: Ordinary objects persist in virtue of being such that, at each time at which the object is present, the fusion of all of the parts of the object that are present at that time is identical to the whole object.

This account attempts to capture the idea that, according to Three-Dimensionalism, at every time, none of the object is missing. That is: though a three-dimensionally persisting object may gain new parts later, there is enough present now to make up the whole object. And that is true of every time at which the object is present: there is enough to make it up. Unfortunately for this account, it classifies paradigm four-dimensional objects as persisting three-dimensionally. For suppose that I am a spacetime worm with proper temporal parts present at each time. At the current time, all of my instantaneous parts are present. However, some extended temporal parts are present as well: they are not contained within the current time, but the current time is not free of them. For instance the

temporal part consisting of all of my life up to and including the current moment is present now. And so is the temporal part consisting of all of my life including and after the present moment. Those two temporal parts fuse to make me. So, in spite of being a spacetime object with temporal parts, I meet the conditions listed for persisting three-dimensionally.

I respond to several accounts (of increasing complexity) that are similar attempts to capture the spirit of Three-Dimensionalism via appeal to Mereological principles. However, each of them fails for similar reasons. We have difficulty capturing the claim that there is enough of the object *contained within each time* to make up the object. The easiest case is the most problematic: consider a three-dimensionally persisting simple, which is intuitively wholly present at each time at which it is present at all, contrasted with a spatiotemporal simple, extended in four dimensions and without smaller parts at any of the proper subregions of the region it fills. These two objects have exactly the same parts (and so have the same mereological features), and if we do not opt for multilocation, they are located at the same region. But one of them persists three-dimensionally, and one does not. It seems, then, that we will need to appeal to something else to set them apart.

Motivated by this case, I turn to attempts to formulate Three-Dimensionalism with appeal to new primitive properties and relations, such as size and shape properties, or distance relations. For instance, at one point Peter van Inwagen suggests that three-dimensionally persisting objects have three-dimensional *size*, though he can remain neutral on where three-dimensionally persisting objects are located. It is important to note that, if we appeal to size, shape, or distance primitives, we cannot understand these in terms of location relations. We cannot, for instance, say that to have three-dimensional size is to be located at a three-dimensional region. If we were to invoke a derivative notion of size of objects such as this one, we would not be able to use size in our formulation of Three-Dimensionalism without thereby invoking location relations and requiring multilocation.

I argue that invoking new primitives, or invoking nothing beyond Mereology, in formulating Three-Dimensionalism does not actually help our three-dimensionalist eternalist to avoid problems caused by my Clifford/Odie case. For the theorist may be able to thereby avoid requiring multilocation, but my cases can actually be stated without using multilocation. All that we require is that Clifford is a proper part of something that is a proper part of it. Typically, we would think we could only find material objects with these features if we have imbedded multilocation. But our three-dimensionalist may have reason to think that some parts can fuse to make an object (which, we might think, grounds it having a particular shape) without the object being located at the region filled by all and only those parts, though the object does have the shape of the region filled by all and only those parts. If this is so, we may think Clifford can be a part of Kibble (because Clifford is a fusion of a collection of objects each of which is a part of Kibble), while Kibble is also part of Clifford (because Clifford is also a fusion of a collection of objects that contains Kibble as a member), even if Clifford is only located at the large region. In this case, Clifford has two shapes (a small one, and a large one), but only one location. The point is this: Three-Dimensionalism in combination with Eternalism requires some sort of *multiple presence*, even if this is not to be cashed out in terms of multilocation. And this sort of multiple presence, however we cash it out, will also be present in my Clifford/Odie case, and thus will give rise to problems for our axioms of Minimal Mereology. We cannot

endorse Three-Dimensionalist Eternalism if we wish to avoid the problems my case produces by restricting what is possible.

8. Implications for Deflationism

In this final chapter, I discuss the implications of these views for deflationism about debates about persistence. Some philosophers think that these metaphysical debates are illusory; when people debate about persistence, for instance, and think they're disagreeing about something substantive, they're really just describing the same thing in two different ways. On this view, there's just matter spread across spacetime, and we're differing in how we use language to divide it up. But: these metaphysical deflationists seem to overlook the fact that their assumption about how matter relates to spacetime is itself a substantive metaphysical claim. And that claim must be argued for. My work can be seen as friendly to these metaphysical deflationists, then, because I begin by assuming their opponent's position (in taking these debates seriously), and end up arguing for a conclusion, namely, that matter is spread through spacetime in a four-dimensional way, which is required for this version of the view that the disagreements are not substantive. Though I do not endorse metaphysical deflationism myself, and my work involves unapologetically and wholeheartedly doing central Metaphysics, my conclusions are informative for even the most Metaphysics-averse.

My hope is that this text will be relatively short, between 150 and 230 pages. The main thesis, that any object's extension across any region requires the object to have proper parts contained within the proper subregions of the region, is modest but very relevant to current literature in Metaphysics. Further, the methodology used to argue for this thesis is helpful and sometimes overlooked, and the thesis has implications for other areas philosophers care about.