

WHERE IN THE RELATIVISTIC WORLD ARE WE?

Cody Gilmore
UC Davis

I do not object to the relativistic reification of “space-time” as a physically real four-dimensional continuum. But I insist that distinct regions of space-time are in some cases “occupied” by numerically the same three-dimensional object. (van Inwagen 1990: 4)

0. Introduction

How do material objects persist through time? The recent literature on this question has centered around two families of views—endurantist and perdurantist.¹ The contrast is often summarized roughly as follows: endurantist views take material objects to lack temporal extent and to persist by being wholly present at each moment of their careers, whereas perdurantist views take material objects to persist by being temporally extended and made up of different temporal parts located at different times.

In this paper I formulate a view about persistence that belongs to the endurantist family. I call the view *saint theory*.² I suspect, but will not try to show, that my formulation of saint theory captures in precise terms what many philosophers have in mind when they use the term ‘endurantism’. More importantly, I believe that saint theory is highly plausible and worthy of detailed discussion in its own right, independently of any relationship it may bear to familiar terms or traditional disputes.

Although there is much to be said for saint theory,³ I do not defend it here. Instead, I raise a problem for the conjunction of this view with orthodox interpretations of relativity theory. Not everyone is wedded to the latter conjunct,⁴ but those who are will see this paper as posing a problem for saint theory itself. Loyal saint theorists, on the other hand, may see it as a reason to doubt orthodox relativity theory. Crudely put, the issue is this: if material objects are in the relevant sense ‘three-dimensional’ and persist by occupying only temporally unextended spacetime regions, how do they ‘fit’ into the four-dimensional,

relativistic spacetimes of Special Relativity (STR) or General Relativity (GTR)? In section 4, I formulate what I take to be the most natural answers to this question and call attention to some serious disadvantages for each of them. Sections 1–3 deal with preliminaries.

1. Spacetime and Exact Occupation

Throughout the paper I presuppose a standard, ‘B-theoretic’ form of spacetime substantivalism.⁵ According to this view, we inhabit a four-dimensional manifold of spacetime points, where points are mereologically simple and unextended, spacetime *regions* are non-empty sets of points, any such set counts as a region, and one region is a subregion of another iff the first is a subset of the second.⁶ The view is substantivalist (as opposed to relationalist) in that it allows for unoccupied regions and makes no attempt to reduce points or regions (or talk of them) to material objects, events, and the spatiotemporal relations they stand in (or talk of them). The view is B-theoretic in that (i) it claims that the past, present, and future and their contents are all equally real, and (ii) it denies that anything instantiates monadic ‘A-properties’ such as pastness, presentness, or futurity. The view is deliberately neutral as to whether spacetime is classical (e.g., neo-Newtonian) or relativistic.

My most distinctive presupposition concerns a relation that can hold between a spatiotemporally located entity (such as a material object, event, immanent universal, or trope) and a region. This relation is sometimes called ‘exact occupation,’ since it is said to hold between a thing and a region just in case, loosely put, the thing *exactly fits into* the region, where this is meant to guarantee that the thing and the region have precisely the same shape, size, and position. Importantly, the relation is intended to be such that there is nothing contradictory or *obviously* impossible about a single thing’s exactly occupying each of two or more regions. I find it highly intuitive, and will presuppose, that we can grasp such relation.

A bit more fully and precisely, I will presuppose that we can grasp a relation of exact occupation that satisfies the following five conditions.

1. It is a necessary truth that if an entity *O* exactly occupies a region *R*, then *O* has, or has-at-*R*⁷, exactly the same shape and size (and geometrical and topological properties more generally) as does *R*, and *O* stands, or stands-at-*R*, in all the same spatiotemporal relations as does *R*.

So, e.g., only a small sphere can exactly occupy a small spherical region.

It is sometimes suggested that a material object can be spherical only in the derivative sense that it can exactly occupy a spherical region. The idea here is that there are actually two different families of shape properties, ‘object shapes’ and ‘region shapes’, and that the former are defined in terms of the latter together

with the exact occupation relation, in the following manner: x is spherical* = df. x exactly occupies a spherical region. *Mutatis mutandis* for other geometrical and topological properties and for spatiotemporal relations. Call this the *Bifurcation Thesis*. I take this thesis to be in conflict with Condition 1. According to that condition, exact occupation holds between a thing and a region only when the (maximally determinate) shape that the *region* has is the very same as the shape that the *object* has; and this is precisely what the Bifurcation Thesis denies.

2. It is not obviously impossible⁸ for an entity to be *multi-located*, i.e., to exactly occupy each of two or more distinct (even nonintersecting) regions.⁹

Perhaps, e.g., there are immanent universals that exactly occupy just those regions that their instances exactly occupy. In particular, suppose that *having a rest mass of 1 unit* is an immanent universal that is instantiated by each electron. Then it will exactly occupy each of the many regions that are exactly occupied by some electron or other. Likewise, perhaps each particular electron has a one-dimensional worldline and is itself multi-located, exactly occupying each point on that worldline.

3. It is not obviously impossible for a multi-located entity to fail to exactly occupy the union of its locations.¹⁰ (Let us say that R is a *location* of O iff O exactly occupies R .)

If there is such an immanent universal as *having a rest mass of 1 unit*, then it exactly occupies a region iff an entity with a rest mass of 1 unit, such as an electron, exactly occupies that region. The union of this universal's locations will intersect any region that intersects an electron's location. Call this (enormous) region R . It is highly unlikely that anything with a rest mass of just one unit exactly occupies R . Since an immanent universal exactly occupies a region only if one of its instances exactly occupies that region, it is also highly unlikely that *having a rest mass of 1 unit* exactly occupies R . But R is the union of that universal's locations.¹¹

4. It is not obviously impossible for a multi-located entity to have certain parts at one of its locations that it lacks at others.

Perhaps, e.g., I exactly occupy both R and R^* and there is a skin cell that is a part of me at R but not at R^* .

5. It is not obviously impossible for an entity, even one without proper parts, to exactly occupy a region R that includes more than one point, while *failing* to exactly occupy any region distinct from R , and in particular, failing to exactly occupy any proper subregion of R .^{12,13}

Perhaps there is an immanent universal *having volume V*, where V is the volume of the entire universe. Then, since there is only one region so large, there is only one region that this universal exactly occupies, even if the universal is simple. Perhaps each electron is a simple but temporally extended entity that exactly occupies only one region—its entire worldline. Perhaps certain fundamental particles are simple and, as a matter of nomic necessity, spatially 1-dimensional in shape. Suppose that such a particle exactly occupies some instantaneous, spatially 1-dimensional region R. The laws of nature guarantee that the particle does not exactly occupy any point-sized subregions of R; and clearly there is nothing incoherent or absurd about the additional assumption that the particle does not exactly occupy *any* proper subregions of R, or indeed any other regions at all.

Although I do not wish to commit myself to a definition of exact occupation, I do not claim that this relation is simple or indefinable. I claim only that it cannot be defined in such a way as to falsify any of 1–5. More could be said, I think, to distinguish this relation from its neighbors. However, for those who will *ever* admit to grasping the relation in question, the list is probably already informative enough to do the job.

Exact occupation, once grasped, can be used to provide straightforward definitions of other interesting relations that can hold between spatiotemporally located entities and regions or points. Consider, for example, the relation of *spatiotemporal (s-) overlap*. Intuitively, a thing s-overlaps a region just in case at least some of that region contains at least some of that thing. The definition in terms of exact occupation would be:

DSO O *s-overlaps* R =df. R is a region, and it intersects some region that O exactly occupies.

Similarly, we can define a relation of a point's *lying within* an entity:

DLW p *lies within* O =df. p belongs to some (point-sized or greater than point-sized) region that O exactly occupies.

A number of philosophers seem to think that there is something suspect or untoward about using a relation of exact occupation (that satisfies 1–5) to define *s-overlap* and *lying within*; these philosophers have expressed a preference for using an undefined relation that behaves like *s-overlap* or *lying within* to define an exact-occupation-like relation (that turns out not to satisfy all of 1–5). Josh Parsons, for example, has proposed a definition structurally analogous to:

DE O exactly occupies R =df. R is a region, and it intersects all and only those regions that O s-overlaps.¹⁴

It is obviously necessary, I take it, that for any object, there is at most one region that intersects all and only those regions that the object s-overlaps.¹⁵ As a result, if exact occupation is defined in terms of DE, it becomes obviously impossible

for a thing to exactly occupy more than one region, and condition 2 is thereby violated.

To my mind, the chief *advantage* of Parsons's approach is its ability to handle the following pair of cases.

Case 1. An instantaneous, unextended point-particle is embedded in a gunky spacetime each of whose subregions is extended. (Something is gunky iff each of its parts has proper parts. Gunky spacetimes do not decompose into simples, be they unextended points or extended 'grains'.) Clearly, our particle does not *exactly occupy* any region in this spacetime; they're all too big. Intuitively, however, we want to say that our particle *s-overlaps* many different regions, where, for any two such regions, there is a third that is a subregion of each of them. If *s-overlap* is primitive, there is nothing to stop us from saying this. If, on the other hand, *s-overlap* is defined by appeal to exact occupation in accordance with DSO, then, since our particle doesn't exactly occupy any regions, it does not *s-overlap* any regions either.¹⁶

Case 2. Suppose that spacetime is 'knuggy,' i.e., that for each spacetime region R, there is some proper superregion of R. (R* is a proper superregion of R iff R is a proper subregion of R*.) Even so, it may seem that there could be an object that, in some sense, is too big to fit into any region smaller than the whole of spacetime. Intuitively, we want to say that this object *s-overlaps* every region. But since there is no *maximal* region (given that spacetime is knuggy), there is no region that our object exactly occupies. Together with DSO, this would entail that the object doesn't *s-overlap* any regions. As before, this problem is avoided if *s-overlap* is primitive.¹⁷

Despite these advantages, approaches like Parsons's do have a drawback: simply put, they deny the quite plausible claim that we can grasp something satisfying 1–5. One response to all of this is to say that *both* exact occupation *and* *s-overlap* are primitive, the principle of ideological parsimony notwithstanding. A second response is to appeal to the three-place relation 'O, at moment t of its proper time, spatiotemporally overlaps spacetime region R,' where a thing can *s-overlap* a region R at one moment of its proper time and fail to do this at another. (Call the relation *s-overlap-at-t*.) We could then offer the following definitions:

DSOS O *s-overlaps* R (*simpliciter*) =df. there is some t such that O *s-overlaps* R at t.

DEO@ O exactly occupies R at t =df. R is a region, and it intersects all and only those regions that O *s-overlaps* at t.

DEOS O exactly occupies R (*simpliciter*) =df. there is some t such that O exactly occupies R at t.

In my view, this proposal is most plausible in the context of a theory that takes moments of proper time to be *sui generis* entities, not reducible in any way to material objects, events, or spacetime points or regions. I suspect that

many philosophers will find the proposal unappealing for this reason. It is worth noting, though, that the proposal's account of s-overlap handles cases 1 and 2 as well as Parsons's approach does, and the proposal's account of exact occupation apparently respects conditions 1–5. I am not aware of any other way to achieve all of these benefits without appealing to more than one primitive occupancy relation.

2. Worms v. Saints

The relation of exact occupation can now be used to formulate saint theory and its opposite, which I call 'worm theory'. Our first step is to define the notion of an entity's *path*. Unofficially, a spacetime region *R* is the path of an entity *O* just in case *R* exactly encompasses *O*'s complete 'career' or 'life-history.' Our official definition will run as follows:

D1 *R* is the *path* of *O* =df. *R* is a region and is the union of the (region or) regions that *O* exactly occupies.¹⁸

Thus, if there is only *one* region that *O* exactly occupies, then *that region* will count as *O*'s path; and if there are *many* regions each of which is exactly occupied by *O*, then the *union* of those regions will count as *O*'s path.

Next we define the notion of an *achronal*, or temporally unextended, spacetime region. Here we take as primitive the relation of absolute chronological precedence, i.e., *being absolutely earlier than*. This notion has a firm basis in both classical and relativistic spacetime theories.

D2 *R* is *achronal* =df. *R* is a region, and for any distinct points *p* and *q* in *R*, neither *p* nor *q* is absolutely earlier than the other.^{19,20}

This definition gives us a notion of temporal unextendedness that meaningfully applies in both relativistic and classical spacetimes. With these last two notions in hand, we can define a notion of *persistence*:

D3 *O* *persists* =df. *O* has a path that is not achronal.²¹

To persist, then, is to have a path that includes a pair of points with absolute temporal order. We are now in a position to formulate a distinction between persisting entities that are singly located and temporally extended (which I call 'worms'), and persisting entities that are multi-located and temporally unextended (which I call 'saints'). See Figs. 1 and 2.

D4 *O* is a *worm* =df. (i) *O* persists, and (ii) for any region *R*, *O* exactly occupies *R* iff *R* is *O*'s path.




Figure 1.

D5 O is a *saint* =df. (i) O persists, and (ii) each region that O exactly occupies is achronal.

Now let me make a few remarks about this contrast.

1. Worm theorists, who hold that all persisting material objects are worms, and saint theorists, who hold that all persisting material objects are saints, can agree about which regions are the paths of which objects. But, for any given persisting material object, worm theorists will disagree with saint theorists as to which regions are exactly occupied by that object.

2. It is important not to interpret saint theory as holding that material objects *move* through spacetime. The claim that things move through spacetime




Figure 2.

seems to entail (i) that there is some temporal dimension, hypertime, distinct from the four dimensions of our ordinary spacetime manifold M , (ii) that the ordinary manifold M exists at different instants, t and t^* , of this hypertime, and (iii) that there are regions R and R^* of M and an object O such that, at hypertime instant t , O exactly occupies region R but not R^* , whereas at hypertime instant t^* , O exactly occupies R^* but not R . Whether or not such a doctrine can be coherently developed, saint theory is in no way committed to it. According to saint theory, there is a ‘tenseless’ relation of exact occupation and each persisting material object ‘tenselessly’ bears this relation to each of many regions.

3. It should be emphasized that even if we confine our attention to persisting entities, the distinction between worms and saints is not exhaustive. From a purely logical point of view, there is room for objects that persist but are neither worms nor saints. Consider, for example, an object that exactly occupies all and only those temporal chunks of its path that have some specified temporal thickness—say, five seconds. Such an object, if it lasts longer than five seconds, would exactly occupy more than one region, and so would not count as a worm; but the regions that it exactly occupies would be non-achronal, and so it would not count as a saint. I do not want to dismiss the possibility of such things or to suggest that they are mere curiosities, but I will not discuss them here.²²

3. (Worms v. Saints) v. (Segmentation v. Anti-segmentation)

The distinction between worms and saints may have some claim to the fiercely contested title, ‘the endurance v. perdurance distinction’. There is, however, a second distinction that may also have some claim to that title. I am tempted to use the familiar phrase ‘temporal parts’ in my characterization of this second distinction, but since that term carries so much unwanted baggage and is itself contested, I will mostly avoid it and use a different term. Very loosely, let us say that x is a *temporal segment* of y iff:

- (i) x ’s path is a ‘temporal slice’ or ‘temporal chunk’ of y ’s path, and
- (ii) x and y ‘share all of their matter’ within x ’s path.

Both clauses can be made more precise,²³ but the task is somewhat tedious and serves no real purpose here, so I will rest content with the current rough definition.

The notion of a temporal segment is meant to be such that if one thing is a ‘temporal part’ of another in the traditional sense, then the first is a temporal segment of the second. But suppose that C is a long-lived lump of clay and that S is a short-lived statue that, throughout its career, shares all of its matter with C . (In that case S ’s path will be a proper ‘temporal chunk’ of C ’s path.) Suppose also that both C and S are saints: they are both temporally unextended objects that persist by exactly occupying each in series of achronal spacetime

regions. Some philosophers may deny that S is a temporal part of C in the traditional sense, on the following grounds: just as only spatially extended things can have spatial parts, only temporally extended things (such as worms) can have genuine temporal parts; hence a *saint* cannot have *temporal parts*. The notion of a temporal segment, however, should be understood in such a way as to guarantee that, given my description of the case, S *does* count as a temporal segment of C, even though both are saints.

Let us say that a persisting thing is *segmented* iff it has a ‘full distribution’ of temporal segments—perhaps a different segment for each instantaneous temporal slice of its path, perhaps a different segment for each continuous temporal chunk of its path, perhaps both. (There is no need to make decision on this issue here.) We will say that a thing is *non-segmented* iff it has no temporal segments at all, aside from itself.

The point that I would like to emphasize here is that, from a purely logical standpoint, the (non-exhaustive) distinction between segmented and non-segmented objects *cross-cuts* the distinction between worms and saints, yielding four more specific types of persisting entities:

- *Segmented worms*. These seem to conform to the picture behind traditional perdurantism, according to which each material object O exactly occupies just one spacetime region R, and each subregion of R (including each temporal chunk and slice thereof) is exactly occupied by a part of O.
- *Nonsegmented worms*. These objects are temporally extended but devoid of anything resembling temporal parts. The most vivid case is that of a simple, persisting point-particle that exactly occupies only one region—its path.
- *Segmented saints*. Consider the relationship, described above, between our long-lived lump of clay C, and our short-lived statue S, both saints. But now suppose that for each subinterval T of the set of times at which C exists, there is a thing x such that (i) x’s career is coextensive with T and (ii) x is made up of the same matter as C throughout T. This makes C a segmented saint.
- *Nonsegmented saints*. These seem to conform to the picture behind traditional endurantism, supplemented with a strict ban on spatial-cum-material coincidence between distinct entities. The most vivid case is that of a simple, persisting point-particle that exactly occupies each point on its one-dimensional path.

For each of these types of entities, there is a proposition that all persisting material objects are of that type. None of these propositions seems to be contradictory or *obviously* impossible, though of course some are more plausible than others. If I am right about this, then, in the general vicinity of the ‘endurance v. perdurance dispute,’ there are actually at least²⁴ two quite distinct and quite separable issues: the worms v. saints issue, and the segmentation v. anti-segmentation issue.²⁵

There may well be entailments between some of the relevant views—e.g., perhaps segmentation theory entails worm theory—but if so the entailments would seem to be non-trivial.

4. The Problem

The remainder of the paper discusses a problem for the conjunction of saint theory with orthodox relativity theory. As will become clear, the problem bears directly on the saints v. worms issue and only indirectly, if at all, on the segmentation v. anti-segmentation issue. The problem arises from the *Location Question*:

What is the general principle that determines, for any given material object, *which subregions of that object's path are exactly occupied by the object?*²⁶

A correct answer to this question will tell us how to fill in the following blank in such a way as to make the resulting principle both true and informative:

For any material object O and spacetime region R, O exactly occupies R if and only if _____.

Unless we are willing to treat the facts about exact occupation as brute and unsystematic,²⁷ we will be attracted to the thought that there must *be* some principle that correctly answers the Location Question. But what is it? Worm theorists and saint theorists will of course disagree.

If each persisting material object is a worm and hence fills its temporally extended path by exactly occupying *only the whole path*, then of course the answer is:

The Path Principle: For any material object O and spacetime region R, O exactly occupies R if and only if R is O's path.

The Path Principle says, in other words, that for any material object there is only *one* spacetime region that is exactly occupied by that object, and the region in question is just the object's complete path. For the worm theorist, then, the Location Question couldn't be any easier. Worm theory immediately *entails* a very simple, definite, and perhaps even tenable answer. Moreover, the worm theorist not only can but must give this same answer in both the classical and the relativistic context.

According to saint theory, on the other hand, persisting material objects trace out their temporally extended paths by exactly occupying *each in a set of achronal sub-regions of those paths*. In the classical context, the saint theorist may seem to have an easy answer. In this context, there is a well-defined equivalence relation of absolute simultaneity, and this relation exhaustively partitions the

four-dimensional spacetime manifold into a set of non-intersecting, global, three-dimensional hyperplanes or ‘time-slices’. Thus, in this setting, it will be natural for the saint theorist to want to adopt:

The Absolute Slice Principle: For any material object *O* and spacetime region *R*, *O* exactly occupies *R* if and only if *R* is an ‘absolute time-slice through *O*’s path’—i.e., iff there is some global hyperplane of absolute simultaneity, *H*, such that $R = \text{the intersection of } H \text{ and } O\text{'s path}$.

To be sure, other saint-theoretic answers to the Location Question are available. Consider, e.g., the Point Principle, according to which each material object exactly occupies all and only the point-sized subregions of its path. The Point Principle is not only consistent with saint theory but entails it.²⁸ However, this principle also entails a very radical and implausible form of spatial multi-location when applied to objects with spatially extended paths. Thus I regard it as a non-starter. In any event, it should be obvious that the Absolute Slice Principle is by far the leading contender in the classical context.²⁹

Once we shift to the relativistic context, however, things become much more complicated and challenging. In this new context, there is no obvious front-runner. Instead there are a number of very different answers that seem to deserve serious consideration. In this paper I discuss what seem to me to be the four most plausible saint-theoretic answers to the Location Question. I dub them the *Every Slice Principle*, the *Rest Frame Principle*, the *Top Down Principle*, and the *Bottom Up Principle*.

4.1 The Every Slice Principle

In the classical context there is exactly one way of dividing up the four-dimensional manifold into a set of non-intersecting, global, three-dimensional time-slices. In other words, classical spacetimes admit of a unique ‘foliation’ into achronal ‘leaves’.

In relativistic contexts, however, absolute simultaneity is not well-defined and, typically, there are *many different* ways of foliating the manifold into global, achronal leaves.³⁰ For example, in the Minkowski spacetime of STR, each different inertial reference frame (or state of unaccelerated motion) picks out a different foliation. Specifically, for any inertial frame *F*, there is exactly one foliation of Minkowski spacetime into the set of global hyperplanes of simultaneity-in-*F*.³¹

This can be made vivid by imagining the straight, four-dimensional path *R* of some everlasting, spatially three-dimensional object *O* that never accelerates. Let *F* be *O*’s rest frame, and suppose that *O* is always spherical, with a radius of five inches, in that frame. Then every hyperplane of simultaneity-in-*F* (i.e., every ‘*F*-plane’) will intersect *R* at a right angle, and each such intersection will be a

spherical achronal region with a radius of five inches. Since, for each point in the entire manifold, there is exactly one F-plane to which that point belongs, the following will also be true: for each point in R, there is exactly one *intersection of R with an F-plane* to which that point belongs. (Let an *F-slice of R* be a region that is the intersection of R with an F-plane.) But now let F* be some distinct inertial frame. Then the F*-planes will intersect R at a different angle than do the F-planes, and each such intersection will be an *oblate* achronal region. Each point in R will belong to exactly one F*-plane and hence to exactly one F*-slice of R. Importantly, each F*-slice of R will intersect many F-slices of R, and vice versa.

Familiarity with this picture, and with the principle that there are no privileged inertial frames, naturally leads to the suggestion that, roughly, each object exactly occupies *each achronal slice through its path*, and only those. So, the thought goes, our everlasting object O described above exactly occupies each of the F-slices through its path, and it *also* exactly occupies each of the F*-slices through its path, and so on.

This rough suggestion can be made precise in significantly different ways, but before we concern ourselves with that task, I would like to present and reply to a pair of considerations in favor of the guiding idea.

First Consideration. It may seem that the Every Slice Principle is just obviously true, given saint theory. Let Omar be a saint, and let R_O be Omar's path. Then, isn't it just obvious that each slice through R_O contains Omar? Each of the relevant slices must contain *something*; each of them is *matter-filled*, after all. And, given that Omar is a saint—an object that traces out its path by exactly occupying each in a set of achronal sub-regions of its path—what could the relevant slices contain aside from Omar itself?

Reply. To see where this line of thought goes wrong, note that there is a distinction between a region's being *matter-filled* and a region's being *exactly occupied* by a material object. We can define the former in terms of the latter: a region R is *matter-filled* iff each subregion of R intersects a region that is exactly occupied by a material object. Thus, e.g., if R is a region each of whose point-sized subregions is exactly occupied by a point-sized material particle, then R will be matter-filled, even if there is no larger object that is composed of the relevant particles and that exactly occupies R. Alternatively, suppose that objects can be extended without having proper parts, and consider a world that contains just one material object, a simple sphere. Further, suppose that this sphere exactly occupies just one region. Each of that region's proper subregions will be matter-filled but none will be exactly occupied by any material object.

As these examples show, it's not at all obvious that being matter-filled entails being exactly occupied by a material object. Returning to the case of Omar and R_O , it's not obvious that each of the slices through R_O is exactly occupied by anything, much less by Omar. For all we know, it may be the case that *only the members of a certain select subset* of these slices are exactly occupied by Omar,

while the rest of the slices are matter-filled but not exactly occupied—not by Omar and not by anything else. Provided that each point in R_O belongs to at least one of these select slices, there's no inconsistency in saying that Omar has R_O as its path without exactly occupying *every* achronal slice through R_O .

Second Consideration. One famous consequence of STR is that the spatial shape of a thing is frame-relative. As I mentioned above, if an object is perfectly spherical in *one* inertial frame, then there will be *another* inertial frame in which the object is oblate, having the shape of a slightly squashed ball. It is also a consequence of STR that each object has some shape in every such frame. But it may seem plausible to suppose that having a shape in a frame F entails exactly occupying an appropriately shaped subregion of some F-plane. This would make it natural to think that, for each inertial frame F, each object exactly occupies each F-slice of its path. Assuming that objects exactly occupy no *other* regions, we arrive at a version of the Every Slice Principle.

Reply. There is a perfectly legitimate sense in which an object can have a shape in a frame F despite the fact that it does not exactly occupy any subregion of any F-plane, viz:

O has shape S in inertial frame F =_{df} (i) S is a shape, (ii) F is an inertial reference frame, and (iii) an F-slice of O's path has S *simpliciter*.

This definition clearly captures *one* sense in which an object can have a shape in a frame; and in this sense, it is not obviously impossible for an object to have a shape in a frame F without exactly occupying any sub-region of any F-plane. For let O be an everlasting object that is always spherical in its rest frame, F, and suppose that O exactly occupies *all and only the F-slices through O's path*. Now let F* be some other inertial frame. Then, since the F*-slices through O's path will be oblate, there will be a sense in which O is oblate in F* despite the fact that O does not exactly occupy any F*-slices of its path. So, unless STR entails that O is oblate in F* *in some stronger sense*, which is doubtful, it seems that we can accommodate the relevant consequences of the theory without admitting that O exactly occupies every slice through its path.

Before turning to criticism of the Every Slice principle, I would like to spend a moment discussing questions about how to formulate it more precisely. So far I have been relying very heavily on the notion of a (flat) hyperplane of simultaneity-in-an-inertial-frame. This notion has application in Minkowski spacetime but not in most General Relativistic spacetimes, which exhibit curvature. A much broader notion, one that has application in both contexts (and even in pre-relativistic contexts), is that of a maximal achronal region. (I take the term from Balashov *forthcoming*, who offers an equivalent definition.)

D6 R is a maximal achronal region = df. (i) R is a spacetime region, (ii) R is achronal, and (iii) R is not a proper subregion of any achronal region.




Figure 3.

Unlike hyperplanes, maximal achronal regions need not be flat; there need not be any inertial frame with respect to which their points are all pairwise simultaneous. Both sorts of regions are meant to be global in the sense of extending across the entire manifold.

Given this distinction between hyperplanes and maximal achronal regions, we can see that there are at least³² two ways of formulating the Every Slice Principle:

The Every Flat Slice Principle (EFS): For any material object O and spacetime region R , O exactly occupies R iff there is some inertial frame F and global hyperplane H of simultaneity-in- F such that $R =$ the intersection of H and O 's path, i.e., iff there is some inertial frame F such that R is an F -slice of O 's path.

The Every Maximal Achronal Slice Principle (EMASL): For any material object O and spacetime region R , O exactly occupies R iff there is some maximal achronal region R^* such that $R =$ the intersection of R^* and O 's path (i.e., such that $R =$ the R^* -slice of O 's path).³³

In the context of a GTR spacetime that possesses no hyperplanes, the 'Flat Slice' formulation, EFS, entails that no material object exactly occupies any spacetime region. This would seem to give saint theorists sufficient reason to prefer the 'Maximal Achronal Slice' formulation, EMASL. I turn now to problems for the Every Slice Principle.

Problem 1: Corner Slices. This is less a *problem* for the Every Slice Principle than an opportunity to become still clearer about how the principle should be formulated. Let M be a molecule composed of ten persisting point-particles, arranged and bonded in the fashion of a chain. Suppose that M inhabits Minkowski spacetime and has a life-history as follows. With respect to M 's rest frame, M and all ten of its particles pop into existence instantaneously. They remain mutually at rest for a few seconds, then, with respect to M 's rest frame, they all pop out of existence instantaneously. Figure 3 depicts M 's path and (parts of) two maximal achronal regions that intersect it.

The lower maximal achronal region, call it R_l , intersects M 's path and the paths of each of M 's ten constituent particles. The intersection of R_l and M 's path is a region that seems to be a suitable container for M : there is nothing implausible about the claim that M exactly occupies that region. Consider, on

the other hand, the upper maximal achronal region—call it R_2 . Although it does intersect M 's path, it intersects the path of only one of M 's constituent particles. Call the intersection of R_2 with M 's path *sub- R_2* . *Sub- R_2* does not seem to be a suitable container for M : it seems too small. And yet both EFS and EMASL entail that M does exactly occupy *sub- R_2* .

Suppose, however, that we adopt a somewhat different formulation of the Every Slice Principle—viz.:

The Every Maximal Achronal Subregion Principle (EMASU): For any material object O and spacetime region R , O exactly occupies R iff R is an achronal subregion of O 's path and is not a proper subregion of any achronal subregion of O 's path.

Presumably, it is possible to find an achronal subregion of M 's path that has *sub- R_2* as a proper subregion. Let F be M 's rest frame. Then, since *sub- R_2* is point-sized, there will be an F -plane that intersects the paths of all ten of M 's constituent particles and that has *sub- R_2* as a proper subregion. The intersection of this F -plane and M 's path is an achronal subregion of M 's path. This shows that, according to EMASU, M does not exactly occupy *sub- R_2* . Perhaps there are other solutions to the problem of corner slices, but shifting to EMASU seems to me the most natural.

Problem 2: Impenetrable Spatially Extended Simples. A material simple is a material object that has no proper parts (parts aside from itself). Perhaps there are no material simples. Or perhaps, while there are some material simples, they are all spatially point-sized, in which case the only achronal regions that they exactly occupy are point-sized regions. Perhaps, however, some material simples are spatially extended and exactly occupy *continuous* achronal regions. I believe that such objects deserve to be taken seriously.³⁴

Suppose, then, that there is a type T of fundamental particle each of whose instances is always spatially extended and continuous. Anyone who is willing to grant this possibility should also grant the possibility that the particles in question, the T -particles, are governed by a law of nature that makes them mutually impenetrable: T -particles cannot pass through one another or have intersecting locations. (A *location* of an object, recall, is a region that the object exactly occupies.)

But this threatens to conflict with the Every Slice Principle. All versions of that principle entail that if a persisting object inhabits a typical relativistic spacetime and is always spatially extended and continuous, it will exactly occupy regions that 'criss-cross' in the manner indicated by Fig. 4. Suppose that the object depicted by Fig. 4 is a T -particle and that it exactly occupies each of the criss-crossing regions. Then isn't it 'passing through itself' in precisely the manner that is forbidden by the law that makes T -particles mutually impenetrable?³⁵ This seems to show that the Every Slice Principle is false in any typical relativistic world containing a group of mutually impenetrable spatially extended simples.



Figure 4.

To the extent that a satisfactory answer to the Location Question must be true at *all* relativistic worlds, this constitutes a drawback for the Every Slice Principle.

Objection. The law barring interpenetration among T-particles is most naturally interpreted as entailing merely that no two *distinct* T-particles can have intersecting locations. The case above involves just a single T-particle whose locations intersect, and this situation does not violate the law. Hence the Every Slice Principle can be true even at worlds containing mutually impenetrable extended simples.

Reply. Consider such a world, and let O be a T-particle within that world. Now suppose that O travels backward in time and has a trajectory that puts it on a collision course with its younger self. Must we suppose that it can freely pass through its younger self, since ‘they’ are not distinct particles? Perhaps there are worlds governed by laws that allow a particle to pass through itself but not through any other particle of the same type. But surely there are worlds governed by laws that do *not* discriminate in this way. Assuming that our object O belongs to one such ‘non-discriminatory’ world, O and its younger self will interact in the same manner as any two T-particles will interact: they will be unable to interpenetrate. It seems to me that, other things being equal, we should leave the possibility of such worlds open. And the Every Slice Principle will fail at these worlds.

Problem 3: Immanent Causation. I find it highly plausible that a material object can exactly occupy two different spacetime regions only if there is an appropriate sort of causal relation (often called ‘immanent causation’) holding between the contents of these regions.³⁶ We can call this the ‘*M*ulti-*l*ocation *R*equires *I*mmanent *C*ausation’ Principle or *MURIC*.

I will not give an account of immanent causation here.³⁷ But the intuitive idea at work in *MURIC* can be indicated very roughly as follows. I am currently in a certain overall intrinsic condition: I have a certain mass, a certain shape, and so on. Call this *condition A*. I was also in a certain overall intrinsic condition at a certain time yesterday. Call this *condition B*. Presumably I am currently in condition A *at least in part because* I was in condition B yesterday. In other words, my being in condition B yesterday *is a cause of* my being in condition A today. So, according to *MURIC*, in order for a material object to exactly occupy distinct regions R and R*, the object’s having the condition that it has in R must be a cause of its having the condition that it has in R*, or *vice versa*. In this way *MURIC* imposes a constraint on *where a thing can be*. A thing cannot exactly occupy each of two regions whose contents do not stand in the appropriate causal relation.

It should be fairly easy to see why MURIC poses a problem for the Every Slice Principle. Let O be a material object that is always spatially extended and continuous, and let R and R^* be achronal slices through O 's path that criss-cross (in the manner of Fig. 4). Suppose that R can be divided into roughly equal-sized subregions, R_1 and R_2 , and that R^* can be divided into roughly equal-sized subregions R_1^* and R_2^* . Since R and R^* are both slices through O 's path, the Every Slice Principle entails that O exactly occupies both R and R^* . See Fig. 5.

But the appropriate immanent causal relation clearly does not hold between the contents of criss-crossing regions such as these. A full half of R —namely R_2 —is such that none of its points lies in the absolute past of any point in R^* . Therefore, no event confined to R_2 causes any event confined to R^* .³⁸ Similarly, a full half of R^* —namely R_2^* —is such that none of its points lies in the absolute past of any point in R . Therefore, nothing that is confined to R_2^* causes anything that is confined to R . This shows, I take it, that the causal relation between the contents of R and the contents of R^* do not count as ‘an appropriate immanent causal relation’. In that case MURIC forces us to conclude that our object O does *not* exactly occupy both R and R^* . This directly contradicts the Every Slice Principle.³⁹

*Objection: Replace MURIC with MURIC**

The foregoing considerations show only that, when we make the shift from a classical to a relativistic context, we need to reformulate MURIC accordingly. A more appropriate formulation of the immanent causal requirement (the objector may suggest) is:

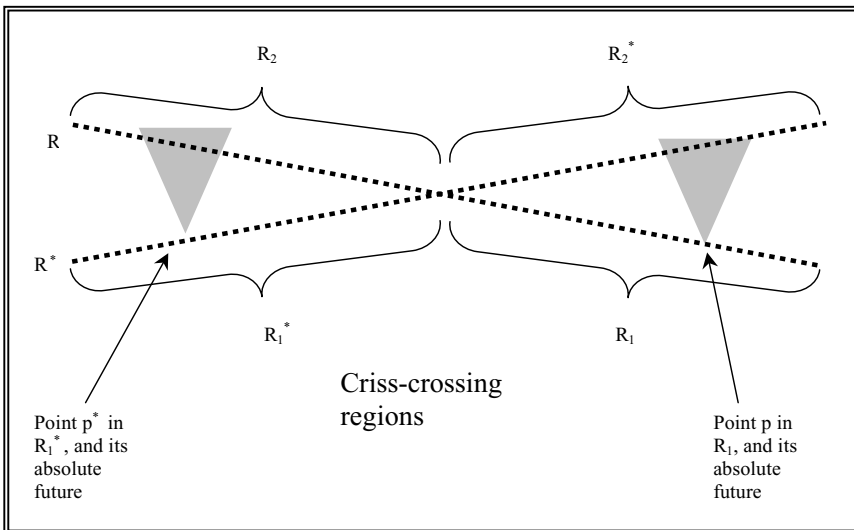


Figure 5.

MURIC* Necessarily, for any material object O and distinct spacetime regions R_1 and R_2 , if O exactly occupies both R_1 and R_2 , then there is some region R such that:

- (i) R_1 and R_2 are achronal slices of R ,
- (ii) there is a set S of achronal slices of R such that every point in R belongs to at least one member of S , and for any two members, x and y , of S , the contents of x bear the appropriate immanent causal relation to the contents of y , or vice versa.

Two points about MURIC* need to be stressed. First, like MURIC, it is an attempt to capture the intuition that material objects cannot ‘jump causal gaps’. Second, unlike MURIC, it poses no problem for the Every Slice Principle. Each of these points deserves some elaboration.

*First point in support of MURIC**. Consider a case of ‘immaculate replacement’.⁴⁰ Let R be a spacetime region that by all outward appearances is the path of a persisting point-particle: R is one-dimensional, timelike, continuous, matter-filled, and is not a proper subregion of any continuous matter-filled region. But now suppose that while the events that lie along the first half of R constitute the appropriate sort of causal chain, and while the events that fill the second half of R also constitute such a chain, there is no causal link whatsoever between the first and second halves of R . Nothing that goes on in the first half of R has any effect on anything that goes on in the second half of R .

Now consider a slice R_1 through the first half of R and a slice R_2 through the second half of R . Is there a single material particle that exactly occupies each of these slices? The intuitively correct answer is No. Despite the spatiotemporally continuous, matter-filled path between R_1 and R_2 , and despite any degree of intrinsic similarity that might hold between their contents (and between the contents of the points on the path from R_1 to R_2), the fact that no causal relation holds between R_1 and R_2 prevents them from both being exactly occupied by the same material object.

To their credit, both MURIC and MURIC* yield this intuitive answer. MURIC says that there is an immanent causal relation, C , that the contents of R_1 must bear to the contents of R_2 (or vice versa) in order for there to be a single material object that exactly occupies both regions. (Perhaps this relation, C , is just the ancestral of the rather demanding, presumably non-transitive *being an appropriately full and appropriately direct immanent cause of* relation, in which case the contents of R_1 bear C to the contents of R_2 just if they stand at opposite ends of some *chain* $\langle e_1, e_2, \dots, e_n \rangle$ such that, for each pair $\langle e_i, e_{i+1} \rangle$ of adjacent links in the chain, e_i bears the more demanding relation to e_{i+1} .) But the requisite relations do not hold. MURIC*, on the other hand, says that in order for a single material object to exactly occupy both R_1 and R_2 , they must both be slices through some path that is, roughly speaking, decomposable into a set of pairwise immanent-causally-related slices. But in the case of R_1 and R_2 , there

is no such path. The best candidate is R; but as we have stipulated, R is *not* decomposable into a set of pairwise immanent-causally-related slices: none of the slices in through the first half of R is immanent-causally-related to any of the slices through the second half of R.

*Second point in support of MURIC**. Unlike MURIC, MURIC* does not bar a single material object from exactly occupying criss-crossing slices, and hence it poses no problem for the Every Slice Principle.

Return to the example of criss-crossing slices depicted in Figure 5. The spatially extended and continuous object O has P as its path, and R and R* are criss-crossing achronal slices of P. Since R and R* criss-cross in the manner indicated, the appropriate sort of immanent causal relation does not hold between their contents. MURIC therefore yields the result that there is no material object that exactly occupies both R and R*.

MURIC*, however, does not yield this result. After all, R and R* are both slices through a path, P, that can be decomposed into the following sort of set: a set S of achronal slices-through-P such that, for each pair of members of S, the contents of these members stand in the appropriate sort of immanent causal relation. We could, for example, arbitrarily pick some inertial reference frame F and divide P into the set of F-slices through P. No F-slice of P criss-crosses with any other F-slice of P. So, if S is the set of F-slices through P, then, for each pair of members of S, it would seem that the contents of these members *are* pair-wise related by the appropriate sort of immanent causation.

*Reply: Why MURIC should not be abandoned in favor of MURIC**

Consider a persisting chain of particles, strung out over a spatially long distance, but such that any two adjacent particles in the chain are bonded together. Let the chain consist of as many particles as you like and be as long as you like. See Fig. 6.

The diagram is meant to show that, *with respect to the rest frame F of the chain*, the following process unfolds: the chain, and all of its constituent particle pop into existence simultaneously. They retain their initial configurations and remain mutually at rest and unaccelerated for a period of time. Then, one by one, starting at the right end of the chain and progressing toward the left, the particles undergo immaculate replacements. The first to go is the particle on the far right: it pops out of existence and is immediately, or almost immediately, replaced with an exact intrinsic duplicate to which it bears no causal




Figure 6.

relation. (Thus the gaps in diagram are meant to represent causal gaps. Whether they represent spatiotemporal gaps depends upon whether the replacements are literally immediate.) Some time later, the same thing happens to its next-door neighbor to the left. And so on down the chain, until each of the originals has been replaced with an exact duplicate.

There is a different inertial frame F^* , however, with respect to which the above process can be described as follows. One by one, some particles pop into existence. They form a chain that is in motion. Then, all at once, the chain and each of its constituent particles undergoes an immaculate replacement.

Let R be the union of the paths of the original particles and the paths of their immaculate replacements. To make things as simple and vivid as possible, suppose that each of the original particle paths has a last point, and let R_1 be the (achronal) region containing all and only these points.⁴¹ Also, suppose that each of the replacement particle paths has a first point, and let R_2 be the (achronal) region containing all and only these points. Then R_1 and R_2 will ‘directly flank the causal gap’. Moreover, nothing that goes on in R_1 will be a cause of anything that goes on in R_2 , or vice versa. These regions will be causally isolated in a very strong sense.

Intuitively, then, it should be clear that no single material object exactly occupies each of them. Now, while MURIC straightforwardly delivers this result, MURIC* does not. To see why not, note that the entire matter-filled region, R , apparently *can* be decomposed into a set of achronal slices whose members are pair-wise immanent-causally related.

Consider the chain’s rest frame, F , and all of the F -slices of R . These slices will intersect the gap, if at all, only at an angle. Consider the F -slice of R that intersects the right-most particle, Righty, at the last moment before its replacement. Call this slice $Rest_1$. Now pick an ever-so-slightly later F -slice, $Rest_2$, and suppose that it intersects the path of Righty’s replacement. But $Rest_2$ will intersect all the other particle paths *before* the corresponding replacements. So there will be a part of $Rest_1$ (its intersection with Righty’s path) and a part of $Rest_2$ (its intersection with the path of Righty’s replacement) whose contents are not causally related at all. Apart from those tiny, minor parts, however, $Rest_1$ and $Rest_2$ (or their contents) are as intimately causally related as any two slices through a spatially extended, persisting thing ever are. If we are committed to tolerating persistence through the gain and loss of particles, it seems that we must not demand causal relations tighter than those between $Rest_1$ and $Rest_2$.

But similar reasoning, it seems, would show that $Rest_2$ bears the needed causal relations to a later slice $Rest_3$, and that $Rest_3$ bears these relations to an even later slice $Rest_4$, and so on, *until we have completely crossed the causal gap*: whereas $Rest_1$ intersects each of the original particle paths and no replacement paths, we will eventually arrive at a slice, $Rest_n$, that intersects each of the replacement paths and none of the originals. Assuming that the relevant immanent causal relation is transitive (as it would be if, as I suggested, it is the ancestral of a more demanding relation), we reach the result that for each pair,

x and y , of F-slices of R , x bears the appropriate immanent causal relation to y , or *vice versa*.

Return now to our slices R_1 and R_2 that flank the causal gap. The foregoing remarks seem to show that R_1 and R_2 are achronal slices through a region, R , that can be decomposed into a set S of achronal slices such that the members of S are pair-wise immanent causally related. In other words, MURIC* does not forbid a material object from exactly occupying both R_1 and R_2 , *despite the extreme causal isolation of these regions from one another*. So MURIC* is too weak to capture the causal constraints on multi-location for material objects. The most natural strengthening is MURIC—which, as we have seen, poses a serious problem for the Every Slice Principle.⁴² Worm theorists, who deny that material objects are multi-located, do not face this problem.⁴³

4.2. The Rest Frame Principle

The main problems for the Every Slice Principle arise from a single feature: this principle entails that spatially extended objects exactly occupy criss-crossing slices through their paths. For those who would like to avoid this entailment, perhaps the most natural answer to the Location Question is the *Rest Frame Principle*.

To state this principle, it will be convenient to make some strong assumptions. Assume that in various General Relativistic spacetimes as well as in Minkowski spacetime, we can define a notion of a ‘rest frame foliation’ that satisfies the following constraints.

1. For any region R , the rest frame foliation of R is an (exhaustive) partition of R into a set S of (nonintersecting) maximal achronal subregions of R .
2. An object’s rest frame foliation carves, or comes as close as possible to carving, O ’s path into slices such that for each slice S , the points in S are all simultaneous-in- F , where F is the rest frame of O at the moment of O ’s career corresponding to S .
3. For any material object O , no matter how large or small O may be and regardless of what sorts of accelerations O undergoes throughout its career, O ’s path⁴⁴ has a unique rest frame foliation.

Visually, we can think of an object’s rest frame slices as intersecting the object’s path at a right angle—or at least as approximating this as closely as possible—and, accordingly, we can call them *orthogonal* (ortho-) *slices*. If an object never accelerates, the ortho-slices through its path will all be parallel to one another. If it does undergo an acceleration, however, its pre-acceleration ortho-slices will be non-parallel to its post-acceleration ortho-slices. The Rest Frame Principle can now be stated as follows.

The Rest Frame Principle: For any material object O and spacetime region R , O exactly occupies R if and only if R is an ortho-slice of O 's path, i.e., iff R belongs to the rest frame foliation of O 's path.

The assumptions that I made in order to state the principle are by no means obviously true. But even if we grant them, the Rest Frame Principle still faces a very serious problem.

Problem: Internal Motion. The problem arises when we consider the case of an object some of whose spatially extended parts are in motion with respect to it—as, for example, one of my red blood cells, Cell, is in motion with respect to me.

Let R be my path and R^* be Cell's path. Presumably the latter is a proper subregion of the former. Now suppose that Cell and I are in motion with respect to each other throughout our careers. According to the Rest Frame Principle, I exactly occupy just the ortho-slices of R . Some of these regions are represented by horizontal lines in Fig. 7. That principle also entails that Cell exactly occupies just those regions that are ortho-slices of *its* path, R^* . Some of these regions are represented by the not-quite-horizontal lines in Fig. 7.

Since Cell and I are in relative motion, the ortho-slices of our paths will not be parallel. As a result, *none of the ortho-slices of Cell's path is a sub-region of any of the ortho-slices of my path.*

According to the Rest Frame Principle, then, none of Cell's locations (none of the regions it exactly occupies) is a subregion of any of my locations. But from this it seems to follow that Cell is never a part of me. After all, it seems to be a necessary truth that, loosely put, if a thing x is, at a certain moment of its career, a part of a thing y at a certain moment of *its* career, then the region that x exactly occupies at the relevant moment must be a subregion of the region that




Figure 7.

y exactly occupies at the relevant moment of *its* career. In slogan form, *parts lie within their wholes*.

This 'Within-ness' principle is a consequence of a stronger principle that is also highly plausible. Loosely put, the stronger principle says that the location of a composite whole is the union of the locations of its proper parts. Both principles entail that if a spacetime point lies within an object, then it also lies within anything that has that object as a proper part. Unlike Within-ness, the stronger principle also entails that if a spacetime point lies within a composite object, then it also lies within some proper part of that object.

How to formulate Within-ness (or the stronger principle) more precisely is a complicated matter. It will depend, for example, upon issues about the reality of the past and future, and upon issues about the nature of the fundamental parthood and occupancy relations.

Suppose, for example, that the following claims are necessarily true: (i) only the present and its contents exist, as presentism asserts, (ii) the fundamental parthood relation is the two-place relation 'x is part of y *simpliciter*', and (iii) the fundamental occupancy relation is the two-place relation 'O exactly occupies spatial region R'. In that case, Within-ness should be formulated as follows:

W(a) Necessarily, if O_1 is part of O_2 , then there are spatial regions R_1 and R_2 such that O_1 exactly occupies R_1 , O_2 exactly occupies R_2 , and R_1 is a subregion of R_2 .⁴⁵

Now suppose that the following are necessary: (i) an eternalist form of spacetime substantivalism is true, (ii) the fundamental parthood relation is the four-place relation 'x, at moment t of its proper time, is part of y, at moment t* of its proper time', (iii) the fundamental occupancy relation is the three-place relation of *exact occupation-at-t*, that is, 'O, at moment t of its proper time, exactly occupies spacetime region R', and (iv) moments of proper time are simple, *sui generis* entities not reducible in any way to material objects, events, or spacetime points/regions. In that case, Within-ness can be formulated thus:

W(z) Necessarily, if O_1 at t_1 is part of O_2 at t_2 , then there are regions R_1 and R_2 such that O_1 exactly occupies R_1 at t_1 , O_2 exactly occupies R_2 at t_2 , and R_1 is a subregion of R_2 .

It should be clear that, in the context of our discussion of the Rest Frame Principle, W(z) is relevant but W(a) is not. It should also be clear that W(z) will, in conjunction with the Rest Frame Principle, entail that if Cell and I inhabit a relativistic spacetime and are always in relative motion, then at no moment of Cell's proper time is it a part of me, at any moment of my proper time. I regard this as a major drawback for the Rest Frame Principle.

Of course, when we change our background assumptions concerning space and time and the fundamental parthood and occupancy relations, we arrive at different formulations of Within-ness. But it seems to me that on any set of background assumptions that are even minimally friendly to saint theory some version of Within-ness will survive and will generate basically the same problem.

It deserves emphasis that the problem depends on Cell's being spatially extended. To see this, suppose that I have a spatially point-sized fundamental particle as a part, and suppose that its path, sub-R, is a continuous, one-dimensional, timelike subregion of my path, R. Then even if the particle is always in motion relative to me, it will still be true that each ortho-slice of sub-R is a mere point-sized spacetime region and, as such, is a subregion of some ortho-slice of my path. So if all of my proper parts are spatially point-sized, the problem disappears completely.⁴⁶ And if Cell always has a complete decomposition into spatially point-sized particles, then it might be the case that all of Cell's *particles* are parts of me throughout Cell's career, even if Cell itself is never a part of me. This might seem to mitigate the difficulty slightly.

Even so, it remains seriously counterintuitive to deny that Cell is ever a part of me, and the Rest Frame Principle remains highly problematic.

4.3. The Top Down Principle

Our search for an answer to the Location Question bears some resemblance to the search for a principled way of defining a physically 'privileged' foliation of spacetime as a whole into global, achronal leaves, each of which can be regarded as an instant of 'cosmic time'.⁴⁷ Perhaps our answer to the Location Question should appeal directly to such a privileged foliation, in the following manner.

The Top Down Principle: For any material object O and spacetime region R, O exactly occupies R if and only if there is a unique, physically privileged foliation of the spacetime to which R belongs into a set S of maximal achronal leaves, and R is the intersection of O's path and one of these leaves (i.e., R is a 'privileged slice' of O's path).

For our purposes it will not be necessary to settle upon any specific definition of 'privileged foliation'. (One prominent proposal is that a foliation is privileged iff it is a 'constant mean curvature' foliation.⁴⁸) Very roughly, the idea is this. In many simplified, idealized models of GTR, spacetime can be foliated so that each leaf orthogonally intersects the path of each material point-particle. Such spacetimes also admit of many other foliations, but each such spacetime has just one 'orthogonal' foliation. This foliation seems to have a special status, since it corresponds (more closely than any other foliation, at least) to the *rest frame of the universe itself*. As applied to a non-idealized spacetime, the idea is that we may be able to identify certain 'major mass points', and foliate the spacetime so that each leaf orthogonally intersects the path of each major mass point. An

example of a major mass point in our universe might be the center of mass of a super-cluster of galaxies.

Although we are not providing a precise definition, we can affirm one quite weak adequacy constraint on any such definition: a spacetime's privileged foliation must be completely fixed by such factors as (i) the spacetime's topological and metric structure, (ii) the distribution of fundamental properties (including any forms of mass and energy that are relevant to GTR) across the spacetime, and (iii) the distribution of fundamental-particle-paths across the spacetime.⁴⁹ No two spacetime that are just alike with regard to (i)—(iii) can differ in how they are carved up by their privileged foliations.

Just as the Rest Frame Principle avoided the main problems facing the Every Slice Principle, the Top Down Principle seems to avoid the problems facing its predecessors. Unlike the Every Slice Principle, it denies that objects exactly occupy criss-crossing slices, and, unlike the Rest Frame Principle, it affirms that if an object has spatially three-dimensional parts, the locations of these parts will be subregions of the object's locations, even if the parts are in motion with respect to the object. The Top Down Principle is vulnerable to new sorts of complaints, however. I will mention two.

Problem 1: Extrinsicness. This criticism is best put forward in the form of a question. Is there any satisfactory definition of 'privileged foliation' that does not have the following consequence?

(C) There could be two four-dimensional regions, R and R^* , such that:

- they belong to the same or different spacetimes,
- both are quite sizeable—as big, e.g., as the region swept out by our solar system, or even our galaxy, over 100 million years,
- they differ only negligibly, if at all, with respect to:
 - internal topological and metric structure,
 - distribution of fundamental properties (e.g., mass and energy), and
 - distribution of fundamental particle paths, but
- R and R^* intersect the leaves of the privileged foliation (of their common spacetime or their respective spacetimes) at very different angles.

Suppose, for example, that R is the path of the Milky Way, and that it runs approximately parallel to the paths of the major mass points in our universe. In that case, it will intersect the leaves of the privileged foliation at right angles. Now suppose that R^* is the path of an intrinsic duplicate (or near duplicate) of the Milky Way that happens to be a 'maverick' galaxy, a galaxy whose state of motion is very different from that of the major mass points in its (perhaps our) universe. In that case, there seems to be some risk that R^* may intersect the leaves of the privileged foliation at an angle that is very far from orthogonal.

Let us suppose that this is so, and hence that the privileged slices of R approximate its ortho-slices, whereas the privileged slices of R^* do not. Suppose that my path, R_{CG} , is a proper subregion of R , lodged somewhere in its middle. Then R^* has a proper subregion, R_{CG}^* , that is a perfect (or near-perfect) duplicate of R_{CG} with respect to topology, metric, distribution of fundamental properties, and distribution of fundamental-particle-paths.

Suppose, however, that the Top Down Principle is correct. Then, since these paths may intersect the privileged leaves at very different angles, the privileged slices of R_{CG} may be very different from those of R_{CG}^* , and as a result the regions that I exactly occupy may be very different from those that my counterpart⁵⁰ in R^* exactly occupies. Perhaps R_{CG}^* intersects the privileged leaves in such a way that its privileged slices are radically ‘length-contracted’, so that their spatial size is much less than that of the privileged slices of R_{CG} . Despite fact that my counterpart and I are molecule-for-molecule duplicates in the sense specified above, and despite the fact that my counterpart’s galaxy is in that sense a molecule-for-molecule duplicate of mine, it might turn out that my counterpart is much shorter than I am!

In this way there seems to be some risk that the Top Down Principle will make a thing’s shape and size depend (non-causally) upon such *radically* extrinsic factors as the distribution of matter and energy in extremely distant parts of the universe. None of the other saint-theoretic principles that we have considered (or will consider) have this consequence; and certainly the worm theorist need not embrace it. The Top Down Principle has this consequence just in case all satisfactory definitions of ‘privileged foliation’ entail (C). I do not know whether all such definitions have this entailment. From the philosopher’s armchair, however, this appears to be a threat.

Problem 2: Non-foliable spacetimes. One famous feature of GTR is that it permits the occurrence of spacetimes that admit of no foliation into global achronal regions whatsoever, much less any privileged foliation. (See Gödel 1949.) As applied to such non-foliable spacetimes, Top Down entails that no material object exactly occupies any region. If one thinks that there are nomically possible worlds in which (i) spacetime is non-foliable and in which (ii) at least one material object exactly occupies at least one region, then one will be forced to concede that Top Down is not nomically necessary. If one also holds that a correct answer to the Location Question must be at least nomically necessary, then one will be forced to reject Top Down. I am inclined to accept each of these steps (on the assumption that orthodox relativity theory is actually true), but they are all highly controversial.⁵¹ Since engaging in this controversy here would take us too far astray, I will say no more about Top Down.

4.4. The Bottom Up Principle

Putting things somewhat metaphorically for convenience we can state the *Bottom Up Principle* as follows.⁵² When a composite object comes into existence,



The locations of the oscillating particle gradually fall out of step with the synchronized locations of the inertial particle. Consider the region R containing just the two end points of the upper-most dotted line. This is a synchrony-slice of the composite object's path, but it is not achronal: its two constituent points are timelike separated. Hence Bottom Up, which tells us that the composite object exactly occupies R , entails that that object is not a *saint*.

Figure 8.

each of its constituent fundamental particles is given its own personal stopwatch that records the elapse of the particle's proper time. At the moment of the composite object's creation, each of these particles exactly occupies some spacetime region or other and each of the stopwatches reads '0 minutes'. The composite object exactly occupies the union of these regions; this is its *initial slice*. We then define a *synchrony-slice* of O 's path as a region R such that, for some real number n , R includes a given spacetime point just in case that point lies within one of O 's constituent particles when that particle's stopwatch indicates that n minutes of the particle's proper time have elapsed. Thus, at O 's synchrony-slices, all of the stopwatches will be synchronized: each of them will display the same number of minutes elapsed.

Now we say that O exactly occupies all and only its synchrony-slices (including its initial slice). Thus O will exactly occupy its 'one-minute' slice (the region made up of those points lying within O 's particles when their stopwatches all show that one minute of proper time has elapsed), O will exactly occupy its '(pi)-minute slice,' O will exactly occupy its 'four-minute' slice, and so on. As for the constituent fundamental particles themselves, we shall suppose that each such particle P exactly occupies all and only the ortho-slices through P 's path.

A number of problems for this account arise from the assumptions that I have made in order to state it. I have been assuming, for example, that (i) every composite object decomposes without remainder into 'fundamental particles' or simples, that (ii) every composite object has a finite past, and worst of all that (iii) composite objects never gain or lose fundamental particles. But even if we somehow managed to state the principle without making these assumptions, grounds for criticism would still remain.

Problem: Parts that fall out of step. Consider an object composed of two simple point-particles, one of which is in inertial motion (or on a geodesic)

throughout its career, and the other of which oscillates rapidly back and forth throughout its career, as indicated in Fig. 8.

STR and GTR both entail that if this situation continues for long enough, eventually we will come to a number n such that the point exactly occupied by the inertial particle when *its* stopwatch reads ' n minutes' will be *timelike*, not *spacelike*, separated from the point that the oscillating particle exactly occupies when *its* stopwatch reads ' n minutes.' Thus the region made up of these two points will fail to be an achronal region, and hence will fail to be the sort of region that a *saint* can exactly occupy. According to the Bottom Up Principle, however, our composite object does exactly occupy the relevant non-achronal region. It appears, therefore, that the saint theorist will be able to give the following argument against the Bottom Up Principle:

- (i) There are nomically possible STR or GTR worlds that contain the relevant sorts of composite objects.
- (ii) No STR or GTR world that contains the relevant sort of composite object is a world at which *both* saint theory and the Bottom Up principle are true.
- (iii) Saint theory is true at all nomically possible STR or GTR worlds.
- (iv) The Bottom Up Principle, like any attempt to answer to the Location Question, is true at all nomically possible worlds if it is true at the actual world.
- (v) Therefore, the Bottom Up Principle is not true at the actual world.

In light of this argument, my guess is that very few saint theorists will find the Bottom Up Principle to be a satisfactory answer to the Location Question.

5. Conclusion

I began the paper by making a series of remarks about the relation of exact occupation. I used this relation to formulate a contrast between *worm theory* and *saint theory*, and I emphasized that the question about whether persisting things are worms or saints is separate from questions about whether things have temporal segments. I then posed the Location Question: Which subregions of a thing's path does the thing exactly occupy? I noted that (i) this question has an easy answer given worm theory and that (ii) given saint theory together with the assumption that spacetime is classical, there is an obvious front-runner, but that (iii) there is a wide range of apparently reasonable answers given saint theory together with the assumption that spacetime is relativistic. Of these apparently reasonable answers, I examined what I took to be the four most plausible, and I argued that each of them has counterintuitive consequences.

In each case, I believe that the problems that I discuss are genuine and deep. To be sure, much more remains to be said: new answers could be suggested, and

each of the answers that I considered could be modified in an attempt to avoid the problems I raised. However I *suspect* that each of the new or modified answers could be subjected to criticisms at least as forceful as the originals.

Some philosophers will embrace the counterintuitive consequences that I have mentioned, and in some cases the debate may be destined to end in stalemate. Some, however, may think, 'Relativity theory itself is counterintuitive, so it should come as no surprise if saint theory has counterintuitive consequences in the relativistic context, nor should this be considered a vice.' In my view, this response is misguided. As far as I am aware, none of the problems that I have mentioned here stems from relativity *per se*; rather, each of them stems from the conjunction of relativity and saint theory, and each can be avoided simply by shifting from saint theory to worm theory. Whether the latter is preferable to the former all things considered is of course too large a question to decide here, but given orthodox relativity, this paper adds to the case against saints.⁵³

Notes

1. See Sider (2001) for an overview.
2. The term 'saint' is meant to conjure up thoughts of multilocation. Some saints of the Christian tradition are said to have been 'wholly present' in two places at once.
3. Some of which I say in my (2004: ch. 6) and, using the term 'endurantism' in place of 'saint theory,' in my (forthcoming).
4. See, e.g., Smith (1993), Tooley (1997), Hinchliff (2000), Craig (2001), Crisp (forthcoming), and especially Monton (forthcoming). Crisp formulates an interpretation of General Relativity that makes it compatible with presentism, the view that only the present and its contents are real. Monton points out that the likely successor to General Relativity, Quantum Gravity, also admits of interpretations that are compatible with presentism.
5. This paragraph merely sketches the relevant views. For a fuller and more careful presentation, see Dainton (2001) and many of the references contained therein.
6. As far as I am aware, none of my main claims depends on the assumption that regions are sets rather than sums or 'pluralities' of points. ('Regions are pluralities of points' is just a loose way of expressing the thought that talk of regions can be paraphrased in terms of plural quantification over, and plural reference to, points.) Hudson (2005: 97) mentions the possibility of treating regions as pluralities but opts to treat them as sums.
7. Some such relativization will be needed if, as seems plausible, both of the following are true: (i) it is possible for a thing to exactly occupy each of two differently shaped regions and (ii) nothing can just plain *have* (i.e., instantiate *simpliciter*) each of two determinate, monadic shapes. I would like to remain neutral, however, as to the precise form that the relativization should take. A number of options seem tenable. For example, if shapes are two-place relations that entities bear to regions, then we can say, 'if O exactly occupies R, then O bears the same shape relation to R that R bears to itself'; and if instantiation (having) can hold between

an entity, a region, and a monadic shape, then we can say, 'if O exactly occupies R, then O has-at-R the same shape that R has (or has-at-R).' *Mutatis mutandis* for other geometrical or topological properties and for spatiotemporal relations. (Note that these options *in no way commit us* to saying that regions are the only fundamental bearers of shapes and that material objects have shapes merely in the derivative sense that they exactly occupy shaped regions. These options allow us to say that the material object and the region it exactly occupies bear the very same (perhaps relativized) instantiation relation to the very same (perhaps relativized) shape. See the next note for more on this.)

It might be objected that this relativizing approach to shapes makes condition 1 too weak to be of any help in characterizing exact occupation. For example, someone might claim that, trivially, if R is any region each of whose points lies within me, then I have-at-R the same shape as R. In fact, however, it is highly implausible, hence far from trivially true, that there is a region at which the Great Pyramid of Giza is approximately spherical, even though it is obvious that there are approximately spherical regions each of whose points lies within the Pyramid. (See my 2003 for further discussion of relativizing techniques, especially in connection with spatiotemporal relations, and for references to the literature.)

8. Here and in the following conditions, I say 'not obviously impossible' rather than 'not impossible' because I do not want assume that it is in fact possible for a thing to, e.g., exactly occupy more than one region. For all I know, the relevant sort of multi-location might turn out to be impossible. But exact occupation shouldn't be defined in such a way as to make this a trivial truth. The relation that I have in mind is such that it is a substantive metaphysical question whether or not it's possible for a thing to exactly occupy more than one region. In my view, this sort of multi-location has a status that many philosophers would attempt to capture by saying that it's at least *conceivable* even if not metaphysically possible.
9. Thus exact occupation cannot be defined as Parsons (forthcoming) defines *exact location*.
10. Barker and Dowe (2003) seem to reject this condition.
11. I give a version of this argument in my (2004). Here I incorporate improvements that I learned from Kris McDaniel in conversation in 2005; I do not know whether he would endorse the present argument.
12. Thus exact occupation cannot be defined as Parsons (forthcoming) defines *being wholly located at* or as Crisp and Smith (forthcoming) define *being wholly present at*.
13. I give a similar characterization of exact occupation in my (2004). Others who appear to recognize the relevant relation include Hudson (2001 and 2005), McDaniel (forthcoming b), and Balashov (forthcoming). Hawthorne (2006: ch. 5) mentions the relation but does not commit himself to its intelligibility.
14. Parsons (forthcoming). His term for the defined relation is 'exact location.' My 's-overlap' corresponds to his 'weak location'.
15. Someone who takes regions to be sets of points might deny this by denying set-theoretical extensionality; someone who takes regions to be sums (i.e., mereological fusions) of points, and who claims that any such sum is a region, might deny this by claiming that it's possible for some points to have more than one sum. Even if one of these bizarre positions were correct, DE would still entail that if

- O exactly occupies R and R^* , then these regions include all the same points and intersect (or overlap) all the same regions. Thus DE would still entail that no object can exactly occupy each of two noncoinciding regions.
16. Parsons (forthcoming) contains a more detailed discussion of this case, which I originally formulated as an objection to Parsons's 'Exactness' axiom. This axiom says (in my terminology) that anything that s-overlaps any region exactly occupies some region. Since Parsons is concerned to defend this axiom as a 'conceptual truth', he sees Case 1 primarily as a *threat* to his system.
 17. Case 2 and the term 'knuggy' are due to Parsons (forthcoming).
 18. If one adopts DSOS–DEOS, a better definition of 'path' is this: R is the path of $O = \text{df. } R$ is a region, and it intersects all and only those regions that O s-overlaps simpliciter.
 19. One problem for this definition arises from spacetimes containing closed timelike curves (CTCs) or 'almost closed' timelike curves. (Roughly, a CTC is a continuous one-dimensional path that (i) is timelike in the sense that it is always directed toward the future with respect to its immediate spacetime neighborhood, and so represents a physically possible trajectory of point-particle with mass, and (ii) is closed in that it forms a loop.) Consider, e.g., a region R containing just two arbitrarily near points p and q that, as we might put it, are 'locally' spacelike related. (Roughly, points are spacelike related iff they are simultaneous in some reference frames, non-simultaneous in others, and not causally connectible). Intuitively, we want to say that R is in the relevant sense temporally unextended. But now suppose that R is embedded within a spacetime that contains closed and almost closed timelike curves, and suppose in particular that there is an almost closed timelike curve running from p to q. Then p and q are timelike related (with p being absolutely earlier than q) and hence R counts as non-achronal. To solve this problem, we need to make precise the intuitive notion of a *locally* achronal region. I take my cue from the following remarks by John Earman:

with CTCs present, *global* Laplacian determinism (which requires a Cauchy surface . . .) is inoperable. But *local* determinism makes perfectly good sense. In any spacetime M, g_{ab} , chronology violating or not, and at any point $p \in M$ one can always choose a small enough neighborhood N of p such that $N, g_{ab|N}$ possesses a Cauchy surface Σ with p [in the absolute future of] Σ . And the relevant data on Σ together with the coupled Einstein-matter equations will uniquely determine the state at p (1995: 171).

(A Cauchy surface is a spacelike hypersurface that intersects every maximal timelike curve exactly once.) It seems to me that if these remarks succeed in capturing the intuitive notion of local determinism, then the following definition ought to capture the intuitive notion of a locally achronal region:

- D2* R is locally achronal = df. (i) there is a spacetime M, g_{ab} , of which R is a subregion, and (ii) there is a region N such that (a) for each point p in R, N is a neighborhood of p, and (b) there is a region R^* that is a Cauchy surface of $N, g_{ab|N}$ and that has R as a subregion.

- If this definition is adequate, then wherever I say ‘achronal’ in the main text, this should be read as ‘locally achronal.’
20. A second problem arises from the ‘Trousers world’ described in Sklar (1974: 306–7). Intuitively, such a spacetime results when we begin with two disconnected 3-spaces (the ‘legs’ of the trousers) that merge at a certain time into a single connected 3-space (the ‘waist’). Let R be an achronal region in the left leg, and let R^* be an achronal region in the right leg. Since there are no lightlike or timelike curves running from any point in R to any point in R^* or vice versa, it seems that the union of R and R^* will count as achronal, according to my definition. But there might be some intuitive sense in which R is located just a few seconds prior to the merger and R^* a great many years prior to the merger, in which case it would seem wrong to count their union as achronal. $D2^*$, suggested in the previous note, seems to avoid this problem as well.
 21. Balashov (forthcoming) objects to D1–D3 on the grounds that they would force the stage theorist to deny that ordinary objects persist. (Stage theorists, such as Sider 2001 and Hawley 2001, identify ordinary objects with what other philosophers would call *instantaneous temporal parts of ordinary objects*, and they say that these instantaneous stages *persist* by bearing some temporal counterpart relation to later instantaneous stages.) Balashov proposes a different system of definitions that does not have this consequence. An alternative solution, and one that does less violence to D1–D3 than does Balashov’s, is to replace D1 with $D1^*$: R is O ’s path = df. R is the union of the (region or) regions that are exactly occupied (either by O or) by any of O ’s temporal counterparts.
 22. Hudson (2001) defends a view in the perdurance family according to which each ordinary, macroscopic persisting thing exactly occupies many overlapping, four-dimensional, nonachronal spacetime regions. Thus he qualifies neither as a worm theorist nor as a saint theorist.
 23. See, e.g., my (2004).
 24. Without committing himself to the intelligibility of exact occupation or the saints v. worms distinction, Hawthorne (2006: ch. 5) formulates several other distinctions in the ‘endurance/perdurance’ neighborhood that cross-cut the ‘segmentation/anti-segmentation’ issue. At the opposite end of the spectrum from Hawthorne are those, e.g. McCall and Lowe (2003), who claim that there are in fact *zero* substantive, not-merely-verbal disputes in this neighborhood.
 25. The thesis that Sider (2001) labels ‘four-dimensionalism’ and defends with his ‘argument from vagueness’ is clearly far more similar to segmentation theory than to worm theory. Heller’s (1984) ‘Body/Body-Minus’ argument is in my view better understood as an argument for worm theory than as an argument for segmentation theory.
 26. Similar questions have been raised by Smart (1972: 6–7) and Rea (1998: 232).
 27. Markosian (1998) responds to van Inwagen’s Special Composition Question by arguing that the facts about composition are brute facts, and McDaniel (forthcoming a) responds to Markosian’s Simple Question by arguing that the facts about which objects are simple are brute facts.
 28. My definition of achronality entails that any point-sized spacetime region (i.e., any region whose only member is a single spacetime point) is achronal.
 29. It is not without its problems, however. Consider a classical spacetime that is ‘cylindrical’ in the sense that it has a closed, circular temporal dimension. Now

suppose that a spatially spherical object O pops into existence, lasts so long that its career winds all the way around the cylinder, coexists with itself for a brief interval, T , then pops out of existence. Suppose also that throughout T , the ‘younger and older versions’ of O are always at a spatial distance from one another. In that case, it seems to me that by far the most plausible description of the situation from the saint-theoretic perspective is that O is spatially bi-located throughout T : for each instant t in T (where an instant is a global hyperplane of absolute simultaneity), there are two spherical subregions R and R^* of t , such that O exactly occupies R and O exactly occupies R^* , but O does not exactly occupy their (non-spherical) union. Unfortunately, the absolute slice principle entails just the opposite: it entails that O exactly occupies neither R nor R^* but only their union. Thus, according to this principle, there is an episode in O ’s life at which it suddenly and discontinuously changes from being spherical and spatially non-scattered to being non-spherical, scattered, and twice as large as it was previously.

30. Some spacetimes permitted by GTR cannot be foliated at all. The most famous is the one described in Gödel (1949). Even in the Gödel spacetime, however, it would seem that the *paths* of most ordinary objects admit of many different foliations into *locally* achronal slices.
31. A global hyperplane of simultaneity-in- F is a non-empty, maximal set of spacetime points that are pairwise simultaneous-in- F . Setting aside spacetimes with abnormal topology, all such planes are achronal.
32. There are of course more than two, given that there are various mathematically expressible conditions we could impose on maximal achronal regions short of flatness.
33. Like my definition of achronality, D2, this principle will need to be reformulated in local terms in order to deliver intuitively correct verdicts as applied to cases involving closed and almost closed timelike curves. If D2* is a satisfactory definition of local achronality, then perhaps we can state a plausible localized version of the Every Slice Principle as:

The Every Maximal Locally Achronal Subregion Principle (EMLAS): For any material object O and spacetime region R , O exactly occupies R iff R is a locally achronal subregion of O ’s path and is not a proper subregion of any locally achronal subregion of O ’s path.

34. See McDaniel (forthcoming a and forthcoming b) for reasons to take such spatially extended simples seriously. Parsons (2000) offers a case in defense of spatially extended simples somewhat differently conceived.
35. It might be objected that (i) according to the Every Slice Principle, spatially 3D objects have locations that intersect only at spatially 2D regions, and (ii) there is something suspicious about laws that prevent spatially 3D objects from having locations that intersect at merely 2D regions (on the grounds that this would prevent certain sorts of contact that should be allowed). In fact, however, every formulation of the Every Slice Principle mentioned in this paper aside from EFS seems to entail that any object that is spatially continuous and three-dimensional throughout a finite period of time will have locations that intersect with each other at 3D regions. To convince oneself of this, it should be sufficient to consider just one example. Take an eternal object in Minkowski spacetime that is always spherical in its rest frame. Now consider a flat slice, S , through this object’s path,

associated with the object's rest frame. Divide S into two hemispheres, S_1 and S_2 . S is just one of many maximal achronal slices through the object's path that have S_1 as a proper subregion. All of our formulations of the Every Slice Principle (aside from EFS) entail that the object exactly occupies each of these slices. But all these slices intersect at S_1 , a 3D region.

36. This principle does not entail that a *universal* cannot exactly occupy a causally unrelated pair of regions. Immanent realists hold that universals are 'wholly present wherever they are instantiated'. (See, e.g., Armstrong 1989, O'Leary-Hawthorne and Cover 1998, and Newman 2002.) One natural interpretation of this doctrine is as follows: Necessarily, for any a universal U and spacetime region R , U exactly occupies R iff there is an entity E such that (i) E exactly occupies R , and (ii) E instantiates U at R (or: E instantiates U at the moment of E 's proper time at which E exactly occupies R). This would seem to commit them to the possibility of a universal that exactly occupies each of two regions whose contents are causally unrelated. This would be the predicament of the immanent universal *having unit negative charge* in a universe that never contains anything but two instantaneous, spacelike separated electrons.
37. Motivation for views in the spirit of MURIC can be found in Armstrong (1980), Shoemaker (1984), and especially Swoyer (1984) and Zimmerman (1997).
38. I am assuming, of course, that there is no superluminal causation and that the relevant spacetime contains no closed or almost-closed timelike curves.
39. Note that the problem does not arise for things that are always spatially point-sized. (Nor do either of the other two problems, of course.) Anyone whose ontology of material objects is confined to such things (certain compositional nihilists, e.g.) is therefore free to adopt the Every Slice Principle.
40. The term is due to Swoyer (1984: 598).
41. If we assume that the immaculate replacements preserve the continuity of the paths, then, given our stipulation that there is a *last* point on each path prior to the replacement, we'll have to deny that there is a *first* point on any of the replacement paths. Since, for vividness, I want R_1 and R_2 to directly flank the gap, I will deny that the replacements preserve continuity. As far as I can see, nothing depends upon this decision. The same points could have been made had I sacrificed vividness to preserve continuity.
42. To be sure, MURIC is not the *only* alternative to MURIC*. Consider, e.g., MURIC**:

Necessarily, for any material object O and distinct spacetime regions R_1 and R_2 , O exactly occupies both R_1 and R_2 only if there is some spacetime region R such that:

- (i) R_1 and R_2 are maximal achronal subregions of R
- (ii) *every* (exhaustive) partition of R into an 'appropriately orderable' set S of maximal achronal subregions of R is such that for each pair, x and y , of distinct members of S , x (or its contents) bears the appropriate immanent causal relation to y (or its contents), or vice versa.

We will say that a sequence S is *appropriately* orderable iff there is a relation of *Precedence* such that (i) for any members R and R^* of S , either R Precedes R^* or vice versa and (ii) for any members R and R^* of S if R Precedes R^* then no point

in R^* is absolutely earlier than any point in R . Now return to the immaculate replacement case involving the chain. $MURIC^{**}$ seems to entail that no material object exactly occupies both R_1 and R_2 the slices flanking the causal gap. After all, for every path to which these regions belong, there is a partition of that path into an appropriately orderable set whose members are not pair-wise immanent causally related. Just take the partition associated with the frame F^* .

But $MURIC^{**}$ seems too strong. First, in the abstract, suppose that the appropriate immanent causal relation holds between some pair of regions. Then it seems that, so far as immanent causation goes, there is no reason why they can't both be locations of the same material object. Why should it matter whether or not some additional, external condition holds, such as that they be slices through a path every appropriate partition of which is such-and-such? Second, a bit more concretely, $MURIC^{**}$ apparently entails that, in our chain case, there is no material object that exactly occupies each of the 'rest frame' (F -) slices through R , even though those slices *do* seem to be related (pair-wise) by the appropriate sort of immanent causal relation. It seems preferable for an immanent causal principle to avoid this entailment as applied to that case. It is not at all obvious to me that a material saint could not persist through such a process.

43. Do similar problems confront traditional perdurantist views? Some advocates of such views have endorsed principles analogous to $MURIC$: some have claimed, e.g., that in order for some instantaneous material objects located at different times to compose an ordinary persisting material object (alternatively: . . . in order for them to compose *anything at all* . . .), the instantaneous objects must stand in an appropriate causal relation. (See Armstrong (1980) and Balashov (2003a and 2003b).) In my view, my argument against replacing $MURIC$ with $MURIC^*$ has much less force when construed as an argument against replacing the perdurantist analog of $MURIC$ with the perdurantist analog of $MURIC^*$. I see no absurdity in the claim that there is an ordinary persisting material object that has one achronal temporal part in $Rest_1$, and a distinct achronal temporal part in $Rest_2$, even if these parts are causally unrelated. What's highly implausible is the claim that the very same material object could exactly occupy both $Rest_1$ and $Rest_2$ given the causal isolation of these regions.
44. Could a single region, R , be the path of two different material objects, O and O^* , whose states of motion somehow differ in such a way that the rest frame foliation of R *with respect to* O differs from the rest frame foliation of R *with respect to* O^* ? Perhaps. In that case we should restate clause 3 as follows: for any <material object O , O 's path R > pair, there is a unique rest frame foliation of R with respect to O .
45. If there are nonspatiotemporal entities (e.g., numbers, properties, propositions, etc.) that can be parts, then all versions of Within-ness should be restricted to spatiotemporal entities.
46. Suppose (i) that, as van Inwagen (1990) argues, the only composite material objects are living organisms, (ii) that no living organism has another living organism as a part, and (iii) that all simple material objects are spatially point-sized. Then my problem for the Rest Frame Principle does not arise. Nor does it arise if there are no composite material objects at all, as compositional nihilists (e.g., Hossack 2000, Dorr 2002) claim.
47. The analogy is far from perfect. First, there is the obvious fact that the Location Question concerns the paths of material objects, not spacetime as whole. Second,

there is the fact that a foliation of spacetime is by definition a division of spacetime into nonintersecting sets, whereas there is nothing contradictory about the claim that saints exactly occupy intersecting regions.

48. See, e.g., Belot and Earman (2001) and Crisp (forthcoming).
49. It is sometimes said that Quantum Mechanics introduces a privileged foliation of its own that may be independent of the foliation that is privileged in the sense discussed in the main text. If this suggestion is correct, then I have no quarrel with the principle that each material object exactly occupies each region that is the intersection of its path and some leaf from the foliation that is privileged *with respect to QM*. But I regard this suggestion as in conflict with orthodox interpretations of relativity theory, and so I ignore it here. For a discussion neighboring issues, see Callender (forthcoming).
50. Given Top Down, it is not at all obvious that I even *have* such a counterpart, for it could well turn out that despite its similarity to R_{CG} , R^*_{CG} is not the path of any persisting material object. Perhaps, in order for some simple particles to compose something, they must stand in a certain multigrade relation, e.g., Life. (See van Inwagen 1990.) And perhaps none of the radically-length-contracted privileged slices of R^*_{CG} contain particles standing in that relation at that region. If so it would seem that R^*_{CG} is not the path of any material object.
51. See Gödel (1949), Savitt (1994), Earman (1995), Yourgrau (1999), Callender (2000 and forthcoming) and Crisp (forthcoming).
52. It is straightforward but tedious to state the principle more literally and precisely. The metaphorical formulation is adequate for present purposes.
53. For helpful comments on this paper, I am grateful to Yuri Balashov, Mark Johnston, David Lewis, Matthew McGrath, Andrew Newman, Oliver Pooley, Jim Pryor, Ted Sider, J. J. C. Smart, Paul Weirich, Jessica Wilson, and audiences at Princeton University, Oxford University, the College of William and Mary, Arizona State University, the University of California at Davis, and the University of Missouri—Columbia.

References

- Armstrong, D. M. 1980. 'Identity through Time', in P. van Inwagen (ed.), *Time and Cause: Essays Presented to Richard Taylor*, Dordrecht: D. Reidel, pp. 67–78.
- Armstrong, D. M. 1989. *Universals: An Opinionated Introduction*, Boulder: Westview.
- Balashov, Y. V. forthcoming. "'Exdurance' Defined", *Philosophical Studies*.
- Balashov, Y. V. 2003a. 'Restricted Diachronic Composition, Immanent Causality, and Objecthood', *Philosophical Papers* 32: 23–30.
- Balashov, Y. V. 2003b. 'Temporal Parts and Superluminal Motion', *Philosophical Papers* 32: 1–13.
- Barker, S. and P. Dowe, 2003. 'Paradoxes of Multi-Location,' *Analysis* 63: 106–114.
- Belot, G. and J. Earman, 2001. 'Pre-Socratic Quantum Gravity' in C. Callender and N. Huggett, eds., *Philosophy Meets Physics at the Planck Scale*, Cambridge: Cambridge University Press.
- Callender, C. 2000. 'Shedding Light on Time', *Philosophy of Science* 67 (Proceedings): S587–S599.

- Callender, C. forthcoming. 'On Finding 'Real' Time in Quantum Mechanics', in W. L. Craig and Q. Smith, eds., *Absolute Simultaneity*, New York: Oxford University Press.
- Craig, W. L. 2001. *Time and the Metaphysics of Relativity*, Dordrecht: Kluwer.
- Crisp, T. forthcoming. 'Presentism, Eternalism, and Relativity Physics', in W. L. Craig and Q. Smith, eds., *Einstein, Relativity, and Absolute Simultaneity*, London: Routledge.
- Crisp, T. and D. Smith, forthcoming. "'Wholly Present" Defined,' *Philosophy and Phenomenological Research*.
- Dainton, B. F. 2001. *Time and Space*, Montreal: McGill-Queen's University Press.
- Dorr, C. 2002. *The Simplicity of Everything*. Princeton University Dissertation.
- Earman, J. 1995. *Bangs, Crunches, Whimpers, and Shrieks*, Oxford: Oxford University Press.
- Gilmore, C. S. 2003. 'In Defence of Spatially Related Universals,' *Australasian Journal of Philosophy* 81: 420–428.
- Gilmore, C. S. 2004. *Material Objects: Metaphysical Issues*. Princeton University Dissertation.
- Gilmore, C. S. forthcoming 'Time Travel, Coinciding Objects, and Persistence,' *Oxford Studies in Metaphysics*, vol. 3.
- Gödel, K. 1949. 'A Remark about the Relationship between Relativity Theory and Idealistic Philosophy', in P. A. Schilpp (ed.), *Albert Einstein: Philosopher-Scientist*, La Salle, IL: Open Court, pp. 557–562. Reprinted in Yourgrau 1990: 261–265.
- Hawley, K. 2001. *How Things Persist*. Oxford: Clarendon Press.
- Hawthorne, J. 2006. *Metaphysical Essays*, Oxford: Oxford University Press.
- Hinchliff, M. 2000. 'A Defense of Presentism in a Relativistic Setting', *Philosophy of Science* 67: S575–S586.
- Hossack, K. 2000. 'Plurals and Complexes', *British Journal for the Philosophy of Science* 51: 411–43.
- Hudson, H. 2001. *A Materialist Metaphysics of the Human Person*, Ithaca: Cornell.
- Hudson, H. 2005. *The Metaphysics of Hyperspace*, Oxford: Oxford University Press.
- Markosian, N. 1998. 'Brutal Composition', *Philosophical Studies* 92: 211–249.
- McCall, S. and E. J. Lowe 2003. '3D/4D Equivalence, the Twins Paradox, and Absolute Time', *Analysis* 63.2: 114–123.
- McDaniel, K. forthcoming a 'Brutal Simples,' *Oxford Studies in Metaphysics*, vol. 3.
- McDaniel, K. forthcoming b 'Extended Simples,' *Philosophical Studies*.
- Monton, B. forthcoming. 'Presentism and Quantum Gravity', V. Petkov and A. Wayne, eds., *The Ontology of Spacetime*
- Newman, A. 2002. *The Correspondence Theory of Truth: An Essay on the Metaphysics of Predication*, Cambridge: Cambridge University Press.
- O'Leary-Hawthorne, J. and J. A. Cover 1998. 'A World of Universals,' *Philosophical Studies* 91: 205–219.
- Parsons, J. 2000. 'Must a Four-Dimensionalist Believe in Temporal Parts?' *The Monist* 83, 399–418.
- Parsons, J. forthcoming. 'Theories of Location,' *Oxford Studies in Metaphysics*, vol. 3.
- Rea, M. 1998. 'Temporal Parts Unmotivated', *Philosophical Review* 107, 225–260.
- Savitt, S. F. 1994. 'The Replacement of Time', *Australasian Journal of Philosophy* 72: 463–474.
- Shoemaker, S. *Identity, Cause, and Mind*, Cambridge: Cambridge University Press.
- Sider, T. 2001. *Four Dimensionalism: An Ontology of Persistence and Time*, Oxford: Oxford University Press.
- Sklar, L. 1974. *Space, Time, and Spacetime*, Berkeley: University of California Press.
- Smart, J. J. C. 1972. 'Space-time and Individuals', in R. Rudner and I. Scheffler (eds.), *Logic and Art: Essays in Honor of Nelson Goodman*, New York: Macmillan, pp. 3–20.
- Smith, Q. 1993. *Language and Time*, New York: Oxford University Press.
- Swoyer, C. 1984. 'Causation and Identity', in Peter French, Theodore E. Uehling, Jr., and Howard K. Wettstein, eds., *Midwest Studies in Philosophy* 9: 593–622, Minneapolis: University of Minnesota Press.

Tooley, M. 1997. *Time, Tense, and Causation*, Oxford: Oxford University Press.

van Inwagen, P. 1990. *Material Beings*, Ithaca: Cornell.

Yourgrau, P. 1999. *Gödel Meets Einstein: Time Travel in the Gödel Universe*, Chicago: Open Court.

Yourgrau, P (ed.). 1990. *Demonstratives*, Oxford: Oxford University Press.

Zimmerman, D. 1997. 'Immanent Causation', *Philosophical Perspectives*, 11, *Mind, Causation, and World*: 433–471.