The Structure of Gunk: Adventures in the Ontology of Space

Jeff Sanford Russell

February 16, 2009
Web Version

1 Points and Gunk

Here are two ways space might be (not the only two): (1) Space is "pointy". Every finite region has infinitely many infinitesimal, indivisible parts, called *points*. Points are zero-dimensional atoms of space. In addition to points, there are other kinds of "thin" boundary regions, like surfaces of spheres. Some regions include their boundaries—the *closed* regions—others exclude them—the *open* regions—and others include some bits of boundary and exclude others. Moreover, space includes *unextended* regions whose size is zero. (2) Space is "gunky".[1] Every region contains still smaller regions—there are no spatial atoms. Every region is "thick"—there are no boundary regions. Every region is extended.

Pointy theories of space and space-time—such as Euclidean space or Minkowski space—are the kind that figure in modern physics. A rival tradition, most famously associated in the last century with A. N. Whitehead, instead embraces gunk..[2] On the Whiteheadian view, points, curves and surfaces are not parts of space, but rather abstractions from the true regions.

Dean Zimmerman has been a relentless source of suggestions and criticism throughout the writing of this paper. Thanks are also due to Frank Arntzenius, Kenny Easwaran, Peter Forrest, Hilary Greaves, John Hawthorne, and Beatrice Sanford for helpful comments.

A version of this paper appeared in *Oxford Studies in Metaphysics*. Ed. by Dean Zimmerman. Vol. 4. Oxford University Press, 2008, pp. 248–278.

[1]David Lewis coined the term "gunk" to describe infinitely divisible atomless objects (Lewis, *Parts of Classes*).

[2]For historical background, see Zimmerman, "Indivisible Parts and Extended Objects".

Three different motivations push philosophers toward gunky space. The first is that the physical space (or space-time) of our universe might be gunky. We posit spatial regions to explain what goes on with physical objects; thus the main reason to believe in point-sized regions is the role they play in our physical theories. So far it looks like points are doing well in that regard despite their uncanniness: all of our most successful theories represent space-time as a manifold of points. But do the points really do important theoretical work, or are they mere formal artifacts, scaffolding to be cast off of our final theory? Modern physics does offer some evidence for the latter. Frank Arntzenius observes that the basic objects of quantum mechanics are not fields in pointy space, but instead equivalence classes of fields up to measure zero differences: point-sized differences between fields are washed out of the theory.[3] This is suggestive: perhaps the points don't belong in the theory in the first place.

A second motivation is interest in possibility rather than actuality. For some metaphysicians, whether or not our own universe's space is pointy or gunky is of only incidental interest; the real question is what ways space *could* be—in the sense of metaphysical rather than epistemic possibility. This question is pressing due to the role of *material* gunk in recent debates. Many philosophers have the intuition that atomless material objects are genuinely possible, even if atomism happens to be true in the actual world, and certain views have been criticized for failing to accomodate this intuition.[4] Furthermore, there are reasons to think that gunky space is the right environment for material gunk. One road to that conclusion is by way of "harmony" principles that link the composition of material objects to the composition of the spatial regions they occupy.[5] Another route to gunky space is by way of contact puzzles.[6] Suppose two blocks are in perfect contact. If space is pointy, then there is a region of space at the boundary between the two blocks. Either this region is empty—which apparently contradicts the claim that the two blocks are touching—or else it is occupied by both blocks—which looks like disturbing co-location—or else some strange asymmetry is at work. Perfect contact may not take place in the actual world, but to the extent that one thinks it a genuine (metaphysical) possibility, this should incline one toward gunk. (It's plausible, for instance, that gunky space is the right environment for corpuscular Newtonian mechanics, where collisions are taken seriously.)

A third reason to investigate gunky space is to formalize the psychology of space. Regardless of whether phenomena like perfect contact are physically or even metaphysically possible, they surely play into "common sense" spatial reasoning. Rigorously capturing this kind of reasoning is important for formal semantics, cognitive

[3] Arntzenius, "Is Quantum Mechanics Pointless?"

[4] E.g. in Sider, "Van Inwagen and the Possibility of Gunk".

[5] C.f. Uzquiano, "Mereological Harmony".

[6] C.f Zimmerman, "Could Extended Objects Be Made Out of Simple Parts?"

science, and artificial intelligence.[7]

These different motivations all call for a clear account of space without points. My interest is in framing theories of gunky space that are as formally adequate and precise as the theories of pointy space on offer—an ambition many share. One pioneer is Alfred Tarski, who presents an elegant formalization of Whiteheadian geometry, which he describes as "a system of geometry destitute of such geometrical figures as points, lines, and surfaces, and admitting as figures only solids—the intuitive correlates of open (or closed) regular sets of three-dimensional Euclidean geometry".[8] (Note, though, that the issues in this paper are independent of space's particular dimension or curvature.)

In addition to framing precise theories, it is useful to construct models of gunky space. The open regular sets are a model for Tarski's geometry, since they satisfy his axioms.[9] This demonstrates that his theory is logically consistent—as long as the Euclidean theory, in whose terms the model is couched, is consistent. (Of course, the model is not meant to show what space really *is*. If gunky space were really made of pointy materials, then the gunky space would only exist if the pointy space did as well—this would make things awfully crowded.)

Recently, though, Whiteheadian space has come under fire: Peter Forrest and Frank Arntzenius have independently offered parallel arguments that challenge the coherence of some gunky accounts.[10] Theories like Tarski's run into serious trouble when we ask questions about the precise *sizes* of regions. In this paper I present a generalization of the Forrest-Arntzenius argument that makes clear which features of space are incompatible. The objection thus sharpened, I take up gunk's defense, considering three responses: Arntzenius's, Forrest's, and a new response. The Forrest-Arntzenius argument requires some concessions, but it does not force us to give up gunky space.

2 Three Kinds of Spatial Structure

Before presenting the main argument in Section 3, I need to lay some formal groundwork for the relevant spatial structures. Along the way I isolate three different constraints on Whiteheadian space.

[7] E.g. Varzi, "Parts, wholes, and part-whole relations".

[8] Tarski, "Foundations of the Geometry of Solids".

[9] An open regular set is an open set which is the interior of its closure. The interior of x ($\mathrm{int}\, x$) is the largest open set contained in x (equivalently, the set of x's *interior points*). The closure of x ($\mathrm{cl}\, x$) is the smallest closed set that contains x (equivalently, the set of x's *limit points*). For example: the open interval $(0, 1)$ is an open regular set; the union of open intervals $(0, 1) \cup (1, 2)$ is not.

[10] Forrest, "How innocent is mereology?"; Arntzenius, "Gunk, Topology, and Measure"; A related argument is given in Hawthorne and Weatherson, "Chopping Up Gunk".

2.1 Mereology

The first of these structures is mereology: a region x may be *part* of a region y (for short, $x \leq y$); synonymously: y contains x, or x is a subregion of y.[11] My main argument only requires two properties of this relation (x, y, and z are arbitrary regions throughout):

$x \leq x$. (Reflexivity)

If $x \leq y$ and $y \leq z$, then $x \leq z$. (Transitivity)

These are very weak conditions: any two-place relation at all has a transitive and reflexive extension.

Some results besides the main argument depend on slightly more structure. Adding one further proposition makes parthood a partial ordering:

If $x \leq y$ and $y \leq x$, then $x = y$. (Antisymmetry)

Throughout this paper I assume these three principles without apology. I don't make use of any mereological principles beyond these without due warning.

I also use the following definitions:

> x and y *overlap* iff there is a region z such that $z \leq x$ and $z \leq y$.
> Regions that do not overlap are *disjoint*.
> x is a *proper part* of y ($x < y$) iff $x \leq y$ and $x \neq y$.

Intuitively, a proper part must have something "left over"—something that "supplements" it. This intuition has formalizations of various strengths:

If $x < y$, then y has a part that is disjoint from x. (Weak Supplementation)

Unless $x \geq y$, y has a part that is disjoint from x. (Strong Supplementation)

An even stronger formulation is that x must have some *remainder* in y—some part which exactly makes up the difference between them.

> A region z is the *remainder* of x in y (or the *mereological difference* of y and x, denoted $y - x$) iff: for every region w, $w \leq z$ iff w is a part of y that is disjoint from x.

A remainder is a *maximal* supplement. This gives rise to the third and strongest supplementation principle:

Unless $x \geq y$, x has a remainder in y. (Remainder Closure)

[11] The definitions and axioms in this section are adapted from Varzi, "Mereology".

Besides taking differences between regions, we can also take sums of regions.

A region f is a *fusion* (or *mereological sum*) of a set of regions S iff: for every region x, x overlaps f iff x overlaps some $y \in S$.

(F)

Intuitively, the fusion of the y's is the region that includes all of the y's and nothing more. If Strong Supplementation holds, then any set of regions has at most one fusion. As we have various supplementation principles, we also have various composition principles. This is one standard principle:

For any x and y, some region $x \vee y$ is a fusion of x and y. (Sum Closure)

A much stronger composition principle is the following:

Any set of regions has a fusion. (Unrestricted Composition)

Under this principle, mereological sums exist for arbitrary finite or infinite collections of regions. All these principles taken together comprise the system of *standard mereology*. Euclidean and Tarskian geometry are both standard, but (to reiterate) I don't assume standard mereology in what follows.

Points have a distinctive mereological feature: they are *atoms*, regions with no proper parts. Pointless space has a corresponding feature: it is atomless.

Every region has a proper part. (Mereological Gunk)

2.2 Topology

Topology describes general shape properties such as connectedness, continuity, and having a boundary. Mathematical orthodoxy casts topological structure in terms of primitively distinguished *open* point-sets. But among the spaces we are concerned with here are those that make no distinction between closed and open regions; so the orthodox approach won't do. Instead I'll follow the tradition of Whitehead, Bowman Clarke, and Peter Röper, which instead characterizes topology in terms of a primitive relation of *connectedness*.[12] Intuitively, two regions x and y are connected ($x \bowtie y$) if they are adjacent, including if they overlap. The corresponding notion in point-set topology is sharing a limit point; i.e. having overlapping closures.[13]

[12] Clarke, "A Calculus of Individuals Based on 'Connection'"; Röper, "Region-Based Topology".

[13] Tarski does not take connection as primitive in his geometry. Instead, Tarski treats *sphericality* as a primitive property of regions. The spheres impose a *metric* structure on the space, which induces a topology. Tarski (with some ingenuity) frames a purely mereological condition for two spheres to be *concentric*. Then we can define "x and y are connected" as "There exists some sphere s such that every sphere concentric with s overlaps both x and y". This parallels sharing a limit point in point-set topology: the set of concentric spheres is an "ersatz point".

With connectedness on hand, we can define a number of other important notions.[14]

> x is a *boundary* of y iff every part of x is connected to y and also to some region disjoint from y.

A boundary region is tightly sandwiched between some region and its complement.

> x is *open* iff no part of x is a boundary of x.
>
> A region x is an *interior part* of y iff $x \leq y$ and x is not connected to any region disjoint from y.[15]

It follows from the definitions that for any $x \leq y$, either x is a boundary of y, or else x contains an interior part of y. In particular, every open region has an interior part. Moreover, except when a region is disconnected from the rest of space, every interior part is a *proper* part.

Whiteheadian space isn't just mereologically distinctive, but also topologically distinctive:

> No region is a boundary. (No Boundaries)

Or equivalently,

> Every region is open. (1)

This entails, and in the presence of Strong Supplementation is equivalent to,

> Every region has an interior part.[16] (2)

I call space that satisfies No Boundaries *topologically gunky*. Topological Gunk is a natural extension of Mereological Gunk: not only does every region have a proper part, it has a part which is strictly inside of it.

[14] The following connection axioms suffice to prove the claims in this section:

> $x \bowtie x$. (Reflexivity)
>
> If $x \bowtie y$, then $y \bowtie x$. (Symmetry)
>
> If $x \bowtie y$ and $y \leq z$, then $x \bowtie z$. (Monotonicity)

[15] When Strong Supplementation holds, the clause "$x \leq y$" is redundant.

[16] (1) entails No Boundaries because if x is a boundary of an open region y, then x is disjoint from y and connected to y at every part, hence not open. (2) entails No Boundaries by the following argument. Suppose x is a boundary of y. *First case:* $x \not\leq y$. By Strong Supplementation, some $z \leq x$ is disjoint from y. But every part of x—and so every part of z—is connected to y, so z has no interior part. *Second case:* $x \leq y$. Every $z \leq x$ is connected to some w disjoint from y, and thus disjoint from x as well—so x has no interior part.

I draw on one further topological notion. A *topological basis* is a collection of open regions with which all of the open regions can be characterized: any open region is a mereological sum of basis regions. An important topological basis in Euclidean space is the set of *rational spheres*: the open spheres with rational center coordinates and rational radii. This basis is *countable*. Any space that is shaped like Euclidean space, whether pointy or gunky, has a countable basis. For instance, the regions in Tarski's geometry that correspond to the rational spheres make up a countable basis there. Since in topologically gunky space every region is open,

> If space is topologically gunky and $\mathcal{B}$ is a topological basis, every region contains an element of $\mathcal{B}$.

2.3 Measure

Besides mereology and topology, regions have sizes. This feature is formalized using a *measure function*, m, which assigns sizes (represented by non-negative real numbers up to infinity) to regions of space.[17] A measure function should satisfy this condition:

> If x and y are disjoint, then $m(x \vee y) = m(x) + m(y)$ (as long as the fusion $x \vee y$ exists). (Finite Additivity)

This generalization of Finite Additivity is an axiom of standard measure theory:

> For any countable set of pairwise-disjoint regions S, if f is a fusion of S then $m(f) = \sum_{x \in S} m(x)$. (Countable Additivity)

With Remainder Closure, this has an important consequence:

> For any countable set of regions S, if f is a fusion of S then $m(f) \leq \sum_{x \in S} m(x)$. (Countable Subadditivity)

The whole is no larger than the sum of its parts. I discuss this condition further in Section 4.2.

Now we turn to measuring gunk. In pointy space, there are *unextended* regions—regions (like points, curves, and surfaces) that have measure zero. Gunky space is supposed to be free of unextended regions:[18]

> Every region has positive measure. (No Zero)

[17] C.f. Munroe, *Introduction to Measure and Integration*.

[18] This condition is proposed by Arntzenius and Hawthorne, "Gunk and Continuous Variation".

Gunky space also shouldn't have discrete "chunks", regions with no parts smaller than some finite size. Rejecting chunks means commiting to the following principle:

Every region has an arbitrarily small part: i.e., for every region x, for any $\varepsilon > 0$ there is some $y \leq x$ such that $m(y) \leq \varepsilon$. (Small Regions)

Small Regions holds automatically in pointy space: every region contains a point with size zero. It also follows from Mereological Gunk, given Remainder Closure and Finite Additivity.[19] The conjunction of No Zero and Small Regions entails (and again given Remainder Closure and Finite Additivity, it is equivalent to):

Every region has a strictly smaller part. (Measure-Theoretic Gunk)

As having an interior part is a natural topological elaboration on having a proper part, having a smaller part is a natural measure-theoretic elaboration. Note that since each of the three ways of being gunky constrains a different sort of structure, they need not all stand or fall together.

3 Disaster

In fact, these structural constraints can't all hold. The following argument by Arntzenius gives a picture of why not.[20] (Forrest offers a parallel argument using a different construction.[21]) For simplicity, consider one-dimensional gunky space—a line—and let the *Big Interval* be a one-inch interval in this space. Let a be the $\frac{1}{4}$-inch segment in the middle of the Big Interval. Then consider the two subregions of the Big Interval flanking a, and let b_1 and b_2 be intervals $\frac{1}{2}$ of $\frac{1}{8}$ inch long (i.e. $\frac{1}{16}$ inch) in their respective middles. This divides r into a, b_1, b_2, and four equal regions surrounding them. In each of the latter four regions, pick out a region that is $\frac{1}{4}$ of $\frac{1}{16}$ inch long (i.e. $\frac{1}{64}$ inch); call these c_1 through c_4. And so on. Call $a, b_1, b_2, \ldots$ the *Cantor Regions*.

The Cantor Regions include one $\frac{1}{4}$-inch interval, two whose lengths add up to $\frac{1}{8}$ inch, four whose lengths add up to $\frac{1}{16}$ inch, etc. So the sum of their lengths is $\frac{1}{4} + \frac{1}{8} + \frac{1}{16} + \cdots = \frac{1}{2}$ inch—strictly less than the length of the Big Interval. And yet: if space is topologically gunky, then the Big Interval is the fusion of the Cantor Regions (I'll defer the details to the main argument below). We have a set of regions

[19] Except in the degenerate case where some infinite region only has infinite subregions. *Proof:* Consider an arbitrary region x with positive finite measure. By Mereological Gunk, x has a proper part y, and by Remainder Closure (it's easy to show) $x = y \vee (x - y)$, so by Finite Additivity $m(x) = m(y) + m(x - y)$. It follows that at least one of $m(y)$ and $m(x - y)$ must be less than or equal to $\frac{1}{2}m(x)$. Repeating this n times procures a part of x no larger than $m(x)/2^n$, which is eventually smaller than any positive ε.

[20] Arntzenius, "Gunk, Topology, and Measure".

[21] Forrest, "How innocent is mereology?"

whose lengths add up to half an inch, but whose fusion is a whole inch long: this is serious trouble.

Before responding to this predicament, I want to sharpen the troubling feeling with my own more general argument.[22]

Theorem. *The following five theses are inconsistent:*

1. *Space has a transitive and reflexive parthood relation.*
2. *Space has a topology with a countable basis.*
3. *Space is topologically gunky.*
4. *Space has a non-trivial countably subadditive measure.*[23]
5. *Every region has an arbitrarily small subregion.*

Proof. Consider a region with positive finite measure—the *Big Region*—and call its measure M. Let the *Insiders* be the basis elements that are parts of the Big Region. Since the basis is countable, there are countably many Insiders, so we can enumerate them. For $i = 1, 2, \dots$, pick a subregion of the ith Insider whose measure is less than $M/2^{i+1}$; call these subregions the *Small Regions*. The Small Regions' sizes add up to strictly less than M.

Even so, the Big Region is a fusion of the set of Small Regions. To show this, we need to show that every region that overlaps a Small Region overlaps the Big Region, and vice versa. The first direction is obvious: each Small Region is part of the Big Region. Now let x be any region that overlaps the Big Region, so x and the Big Region have a part in common, y. By Topological Gunk y contains a basis element b, which must be one of the Insiders, and so b contains a Small Region. Since $b \leq x$, x overlaps a Small Region.

Therefore the Big Region fuses the Small Regions, as advertised. But the Small Regions are too small for that! This means that Countable Subadditivity fails. □

Standard mereology, topology, and measure theory, even when stripped to a bare skeleton, are inconsistent with topological gunk. Any account of gunk that commits to all five theses is incoherent—so things look bad for Whiteheadian space.

[22] My argument is more closely analogous to Forrest's variant than Arntzenius's. Forrest's construction, though, depends specifically on pointy representations of regions: he appeals to regions that are well-represented by neighborhoods of points with rational coordinates. Moreover, both Arntzenius's and Forrest's constructions depend on assumptions about the geometry of space. By contrast, my construction is purely topological.

[23] By "non-trivial" I mean that at least one region has positive finite measure.

3.1 Aside: Measurability

My argument officially requires a measure function that is defined for *every* region of space. This may seem too strict a requirement; the Vitali theorem and the more notorious Banach-Tarski theorem show that in Euclidean space there are sets of points that cannot be consistently assigned a measure (these theorems require the Axiom of Choice). Thus the standard measure function in Euclidean space (Lebesgue measure) is defined only on a distinguished collection of sets, which are called *measurable*. Here are two metaphysical interpretations of the mathematics: one might say that some regions in Euclidean space have no sizes; or one might say that the true regions in Euclidean space correspond to just the measurable sets.[24] (On the latter interpretation, Unrestricted Composition fails: an arbitrary set of regions does not necessarily compose a region—for there are sets of point-sized regions which do not. Fusions of arbitrary *countable* sets of regions still exist, but all bets are off for the *uncountable* sets.)

We hope that gunky space will free us from such discomforts. Arntzenius and Hawthorne propose a requirement on gunk to this effect (they call it "Definition"):[25]

> Every region has a measure. (Measurability)

Note that my argument requires no composition principles, so Measurability can be had cheaply by banishing any supposed exceptions, as in the latter response to Banach-Tarski. But even admitting non-measurable regions gives no relief from my argument. Small Regions guarantees the existence of regions with small measures, which are *a fortiori* measurable, and this supplies enough measurable regions to make the argument work. Even if we revise this thesis to say only that every *measurable* region has an arbitrarily small part (so we don't get measurable regions for free) the argument still succeeds with an unobjectionable extra premise: every basis element (e.g., every sphere) is measurable. This permits restricting the measure function far beyond what would suffice to accommodate Banach-Tarski, without escaping an inch from my argument. Measurability is not the culprit.

4 Escape

The generalized Forrest-Arntzenius argument doesn't leave very many places for the gunk theorist to turn. One way out is to reject one of the three kinds of structure: in fact space does not have, or (for the theorist who only wants the *possibility* of gunk) does not necessarily have measure, topology, or mereology. Deleting any of these is dizzying; perhaps sizeless or shapeless spaces are broadly possible even so, but henceforth I

[24] Arntzenius, "Gunk, Topology, and Measure".

[25] Arntzenius and Hawthorne, "Gunk and Continuous Variation".

assume that our gunk theorist wants to hold onto at least the rudiments of parthood, size, and shape. This theorist has to deny one of the five inconsistent conditions on these structures.

The mereological condition—that parthood is reflexive and transitive—is so weak as to be beyond reproach.

What about giving up a countable topology? There are topological spaces that do not have countable bases, but generally speaking they are exotic infinite dimensional affairs. Such space would be shaped nothing like Euclidean space or any other ordinary manifold. There may still be interesting possible spaces to be explored in this direction, but for now I leave that exploration to others.

And Small Regions? We could imagine "chunky" space with some smallest volume—say on the Planck scale—which is a floor on the size of regions. But for space like this to be *gunky* it would have to be non-standard—for as I observed in Section 2.3, denying Small Regions while holding Mereological Gunk entails that either Remainder Closure or Finite Additivity fails. In what follows I consider some revisions to standard mereology and measure theory in their own right, but none of them gain appeal for the gunk theorist by being conjoined with the denial of Small Regions. So I pass over this route, too.

This leaves two claims open to scrutiny: Topological Gunk and Countable Subadditivity. Rejecting either of these bears a serious cost, but does present live options for the gunk theorist. Arntzenius's response denies Topological Gunk, while Forrest's response (in a roundabout way involving non-standard mereology) denies Countable Subadditivity. I will discuss each of these avenues, as well as a third response that appeals more directly to non-standard measure theory.

4.1 Gunk with Boundaries

Arntzenius (following Roman Sikorski and Brian Skyrms) proposes a conception of gunky space that departs rather dramatically from Whitehead's.[26] Arntzenius's space obeys standard mereology, it has a topology with a countable basis, it has a countably additive measure function, and moreover, it is both mereologically and measure-theoretically gunky. It is not, however, topologically gunky, and thus escapes the problem at hand.

To motivate Arntzenius's construction of a model for this kind of space, let's revisit Tarski's model—the open regular sets—from another angle. We start with orthodox pointy Euclidean space and then "smudge" it, leaving out some distinctions the pointy theory makes but retaining the most important structure. There

[26] Arntzenius, "Gunk, Topology, and Measure"; Sikorski, *Boolean Algebras*; Skyrms, "Logical Atoms and Combinatorial Possibility".

is a general method for "smudging" certain unimportant elements out of a mereology—Arntzenius's metaphor for this is "putting on blurry glasses".[27] The method works as long as the unimportant elements form an *ideal*: any part of an unimportant element must be unimportant, and the fusion of any two unimportant elements must be unimportant. To model Tarski's space, we smudge the boundaries of regions; it's not hard to check that the boundary regions are an ideal.

Once we have picked an ideal of unimportant elements, we define an equivalence relation on the original space: two regions x and y are considered equivalent just in case $x - y$ and $y - x$ are both unimportant.[29] (It is easy to show that for any ideal this is indeed an equivalence relation.) This presents us with a natural collection of "smudged" objects: the equivalence classes under this relation represent regions of gunky space.[30] In the case of Tarskian space, this amounts to considering two Euclidean regions equivalent if they agree up to their boundary points.

The representatives inherit mereology from the original space: if gunky regions r and s are represented by equivalence classes X and Y, then $r \leq s$ iff for every $x \in X$ and $y \in Y$, $x - y$ is unimportant. They also inherit topological structure from the Euclidean topology: $r \bowtie s$ iff for every $x \in X$ and $y \in Y$, $x \bowtie y$. Showing that these satisfy the standard axioms is straightforward.[31] (Choosing a measure for Tarskian space is more involved; I discuss this in Section 4.4.) Each Tarskian equivalence class contains a unique open regular set, which we can treat as a region's canonical representative; in fact, the model we have just constructed is isomorphic to Tarski's.

[27] Assuming Sum and Remainder Closure—i.e., that the mereology forms a Boolean algebra. The result is called a *quotient algebra*.[28]

Note: a Boolean algebra conventionally includes "bottom" and "top" elements: a *null region* contained by everything and a *universal region* containing everything. There are trivial translations between mereology with a null region and without it: for instance, in the presence of a null region, disjoint regions have exactly *one* common part. Introducing a null region is a technical convenience on par with introducing ∞ as a value for functions that increase without bound, or # to indicate truth-value gaps.

[29] Including when $x - y$ (or $y - x$) is null—i.e., when $x \leq y$ (or $y \leq x$).

[30] One of these equivalence classes—the ideal itself—represents the null region. This class can simply be discarded from the final structure.

[31] More carefully: in general this process results in a Boolean algebra; whether the algebra is *complete*—i.e., whether Unrestricted Composition holds—needs to be checked in individual cases. In the instances I consider, it does hold.

It's not clear what the "standard axioms" are for connection, but at any rate the defined relation is reflexive, symmetric, and monotonic (Note 14), as well as obeying Röper's decomposition principle,

If $x \bowtie (y \vee z)$ then $x \bowtie y$ or $x \bowtie z$.

Finding a full adequate set of topological axioms in terms of connection remains open. Röper offers a set of ten axioms, but his tenth axiom is equivalent to the problematic No Boundaries condition (and thus is not satisfied by Arntzenius's space); removing this axiom undoes his method of recovering points from pointless topology.

This model gives primary importance to the topological gunk condition: we use blurry glasses that can't see boundaries, ensuring that the No Boundaries condition will be met. As we have seen, this condition is problematic. Arntzenius responds by instead privileging the measure-theoretic features of gunk, in particular No Zero. His model is constructed using blurry glasses that are blind to Euclidean regions with (Lebesgue) measure zero. These regions also form an ideal, so we can follow exactly the same construction: gunky regions are represented by equivalence classes of regions that have measure zero differences. We obtain mereology and topology as in the preceding paragraph,[32] and we can set the size of a region to be the (Lebesgue) measure of any member of its representative equivalence class, since equivalent regions have the same measure. This yields a model of measure-theoretically gunky space couched in the terms of standard pointy space. It follows that the consistency of the one reduces to the consistency of the other: the prospects look bright for Arntzenian space.

On the other hand, this space has some surprises for the gunk theorist. Consider again the Cantor Regions from Section 3. Arntzenian space harbors both the Big Interval and a proper part of it that fuses the Cantor Regions—the *Cantor Fusion*. Moreover, it has a third region making up the difference between them—the *Cantor Dust*. This is a peculiar region. While it is extended (its size is positive), it has no interior, and it is a boundary of the Cantor Fusion: true to form, Arntzenian space violates the topological gunk condition. Furthermore, the Cantor Dust is entirely scattered: every part is divisible into disconnected parts. It is tempting to describe the region as a fusion of uncountably many scattered points—but of course this is the wrong description, since there are no points.

In Euclidean space, one can completely characterize a region in terms of the points that make it up: the points are a *mereological basis*. Likewise in Tarskian space, the rational spheres are a mereological basis. Arntzenian space has no analog to either the points or the spheres in this respect. The points are an *atomic* basis for Euclidean space; the spheres are a *countable* basis for Tarskian space. If Arntzenian space has any set of "fundamental parts" it must be neither atomic nor countable. This isn't a crippling deficit, but it helps explain why regions like the Cantor Dust are difficult to get an intuitive grip on.

[32] Arntzenius uses a different definition of connection. Say p is a *robust limit point* of x iff for every open neighborhood $u \ni p, m(u \cap x) > 0$. (Measure-equivalent sets have the same robust limit points; we can treat the "robust closed sets" as canonical representatives.) Arntzenius's condition for the connectedness of regions represented by X and Y is that some p is a robust limit point of every $x \in X$ and $y \in Y$. This is equivalent to my condition. (Say x and y have no robust limit points in common. Then every point in $\mathrm{cl}\, x \cap \mathrm{cl}\, y$ has a neighborhood u such that $m(u \cap x) = 0$ or $m(u \cap y) = 0$. Since the Euclidean topology has a countable basis, there are countable collections of open sets $\{v_i\}$ and $\{w_i\}$ such that $\mathrm{cl}\, x \cap \mathrm{cl}\, y \subset \bigcup v_i \cup \bigcup w_i$ and for each i, $m(v_i \cap x) = m(w_i \cap y) = 0$. By countable additivity, $x - \bigcup v_i$ and $y - \bigcup w_i$ are measure-equivalent to x and y (respectively) and they have disjoint closures—that is, they are not connected. The converse is clear.)

Arntzenian space is gunky, but in a weaker sense than full-blooded Whiteheadian space. Space could be gunky in a weaker sense still, satisfying Mereological Gunk but neither the topological nor the measure condition. The now-familiar smudging method produces a model for such space: we look at Euclidean space through blurry glasses blind to regions made of *finitely many points*, and define mereology, topology, and measure in parallel fashion to Arntzenius's space. The resulting space underscores how separate the different gunk conditions are from one another: this space is very much like ordinary Euclidean space in that it has boundaries, unextended regions, zero-dimensional regions, distinctions between open and closed spheres—but even so, it has no *atomic* regions—no points. For those whose interest in gunk springs from purely mereological concerns, this "bare-bones" gunk should suffice.

4.2 Interlude: Foundations of Measure

One way to save gunk is to give up the topological condition and settle for a weaker kind of gunkiness. But some gunk theorists—for instance, those who turn to gunk to solve contact puzzles—may be less willing to admit violations of No Boundaries. For these theorists, the main alternative is to deny Countable Subadditivity. This is a natural response for those in the Whitehead-Tarski tradition: the Forrest-Arntzenius problem arises because Whitehead and Tarski simply didn't take measure into account in their theories. In particular, Tarski's "geometry of solids" satisfies four of the contradictory five theses, but breathes not a word of measure.

Since in standard mereology Countable Subadditivity follows from Countable Additivity, I turn to the latter principle. Is there a plausible theory of full-fledged measurable space without this condition? Answering this question requires a detour into the philosophical foundations of measure theory. So far I have treated sizes as if they were fundamental numerical tags on things. This is obviously a fiction; the numbers and arithmetic operations on them arise from other features of the world. Treating sizes as numbers is convenient, but can mislead: it gives the illusion that numerical operations are automatically meaningful, obscuring underlying assumptions.[33] We must make sense of sums of sizes, and specifically infinite sums.

Say we have some determinable property of regions. For this property to be *size*, it has to have quantitative structure. First, the determinates should at least be ordered: it must be sensible to say that this object's size is greater or less or the same as that object's size. This guarantees only a fairly weak kind of quantitativeness—in the medieval jargon, that of *intensive qualities*. Size, though, is an *extensive quality*: it has algebraic structure such that sizes can be *added*. This operation isn't literally numerical addition, but it's structurally analogous. (I'll denote it $\oplus$.)

[33] C.f. Nagel, "Measurement", p. 313.

The algebra of sizes does not float free: it must correspond to a natural relation on the objects that have sizes; I'll call this relation *aggregation*.[34] Aggregation imposes an algebraic structure on things that have sizes; this structure is what underpins addition of sizes. The sum of two things' sizes is the size of their aggregate.

> If x is the aggregate of y and z, then the size of x is the sum of the size of y and the size of z (i.e., $m(x) = m(y) \oplus m(z)$).
>
> (Finite Additivity)

This additivity condition, I claim, has the status of a conceptual truth: without it there is no intelligible talk of adding sizes at all.

What is aggregation? It's natural to say it's mereological structure: x is the aggregate of y and z just in case y and z are disjoint regions whose mereological sum is x. Then the Finite Additivity condition above amounts to the condition by that name in Section 2.3: the sum of two sizes is the size of the mereological sum of disjoint regions that have those sizes. Composition is what gives additive structure to sizes. For now, though, I'll go on saying "aggregation" without committing to which relation it is.

This accounts for finite sums of sizes; but before we can evaluate Countable Additivity, we need to make sense of infinite sums. One way would be to define infinite sums of sizes directly in terms of another primitive relation: aggregation of infinite collections of regions. This would put Countable Additivity on par with Finite Additivity as a defining condition for sums of sizes. But there's another way of spelling out infinite sums of sizes, which is the standard way of making sense of infinite sums of *numbers*. This way leverages the account we already have of finite sums: the sum of an infinite set of sizes S is the least upper bound (if there is one) of all of the finite sums of sizes in S; i.e., $\text{lub}\{s_1 \oplus \cdots \oplus s_n : s_1, \ldots, s_n \in S\}$.[35] Using numbers to represent sizes and operating on them with the $\sum$ symbol tacitly calls on precisely this definition (or something equivalent); so it was implicit in the original statement of Countable Additivity.

This definition of an infinite sum of sizes makes no appeal at all to *aggregation* of infinitely many things. Still, there is such an infinite aggregation relation—for mereological sums, we have this in Definition F. This raises a substantive question: is the size of an aggregate of a countable infinity of regions always equal to the infinite sum

[34] Why think this? One reason is *operationalism*, the view that physical quantities like size must be defined in terms of measurement procedures, and operations like addition of sizes must likewise reduce to some mechanical procedure that can be performed on the things that are measured. Another reason is *nominalism*: if size properties aren't reified, then their algebraic structure must latch directly onto the things that bear the properties, as there isn't anything else for that structure to describe. I don't find either of those reasons quite compelling, but even so I think the principle is true.

[35] "Least upper bound" is understood in terms of the ordering of sizes: $\text{lub}\, S$ is the smallest size that exceeds every size in S.

of their sizes? That is, is the following true?

If x is the aggregate of a countable set of regions S, then the size of x is the sum of the sizes of the members of S. (Countable Additivity)

This principle could go either way. Infinite sums of sizes are defined apart from infinite aggregations of regions; there is no conceptual reason why the two need coincide. Unlike Finite Additivity, Countable Additivity does not spring from the requirements for extensive qualities: it might fail and yet leave size structure intact.

The present escape route involves denying Countable Additivity in the particular case where aggregation is the disjoint fusion relation. One could deny this by rejecting additivity—the limit of a sequence of sizes might not reliably yield the size of an infinite aggregate—or by rejecting fusions—the mereological sum might not be the right relation to play the aggregation role for sizes. I'll consider each of these options.

4.3 Non-Standard Mereology

Forrest levels his version of the argument against standard mereology; in particular, Weak Supplementation.[36] This move appears to miss the point: my strengthened argument does not depend on any supplementation principle, so how could rejecting Weak Supplementation constitute a response? Only if this denial is followed up by another: denying that fusions have countably additive sizes. If Weak Supplementation fails, then the definition of "fusion" (Definition F) pulls apart in many ways from its usual theoretical role—and so in particular it becomes plausible that the fusion relation so defined does not play the *aggregation* role for measure theory.

When Weak Supplementation fails, some sets of regions have more than one fusion; a region may fuse a set and yet fail to contain any of its members; indeed, the fusion may be a proper part of some member. These bizarre consequences suggest that in the absence of Weak Supplementation, the standard definition of fusion fails to capture a natural or theoretically important relation, and fails to capture the ordinary intuitive concept of *composition*.[37] This is not terribly surprising, seeing as Definition F was framed with standard mereology in mind.

Let's distinguish "fusion" from another, closely related concept:

A region u is the *least upper bound* of S (for short: $\operatorname{lub} S$) iff: for every region x, $x \geq u$ iff x contains every $y \in S$.

[36] Forrest, "Nonclassical Mereology and Its Application to Sets".

[37] C.f. Forrest, "Mereological summation and the question of unique fusion". Kit Fine has argued for a general conclusion along the same lines.

(Antisymmetry guarantees uniqueness.) The least upper bound of S is the smallest region that contains every element of S. If Remainder Closure holds, then x is a fusion of S if and only if x is the least upper bound of S. So in standard mereological contexts, the two concepts have the same extension. This fact, I submit, is the main motivation for the standard definition of "fusion". But when supplementation principles fail, these two concepts pull apart from one another, and so we lose the main reason for accepting the definition. In particular, we lose our reason for thinking that the definition captures the right relation to ground our measure theory.

Instead we should turn to the least upper bound. This relation gets along without Weak Supplementation better than the fusion relation: a least upper bound is unique and contains its members, at least. This suggests that we should replace our original statement of Countable Additivity with this one:

$$\text{For any countable set of pairwise-disjoint regions } S,\ m(\operatorname{lub} S) = \sum_{x \in S} m(x).$$

This revision makes room for a consistent theory of topologically gunky space—at the price of adopting (in Forrest's words) a "semi-standard" mereology.

Forrest's model for such space is the open sets with measure zero differences "smudged" (i.e. measure-zero equivalence classes of open sets). The ontology of Forrest's space is richer than Tarski's but poorer than Arntzenius's. We can see this by considering the Cantor Regions from Section 3 once again. Arntzenius's space includes the Big Interval, the Cantor Regions, the Cantor Fusion (the sum of the Cantor Regions), and the Cantor Dust (the difference between the interval and the Cantor Fusion). Tarski's space, meanwhile, includes the Big Interval and the Cantor Regions, but countenances neither the Cantor Fusion (as an entity distinct from the Big Interval) nor the Cantor Dust. Forrest's space includes the Big Interval, the Cantor Regions, and the Cantor Fusion (better called "the Cantor Lub")—but there is no "dusty" region making up the difference between the Cantor Fusion and the interval. This might strike some as the right balance, since the Cantor Dust is the most bizarre of the regions under consideration. But accepting this ontology sacrifices Weak Supplementation: without the Dust, the interval has a proper part with nothing left over. If you paint each of the Cantor Regions in Forrest's space, you thereby succeed in painting the Cantor Fusion but fail to paint the Big Interval—and there's no way to finish the job without repeating some of your work, since every part of the Big Interval has paint on it. This is an unnecessarily drastic concession.

4.4 Finitely Additive Gunk

Returning to more hospitable mereology, why think that additivity holds for countably infinite fusions? One might motivate this with an infinite "supertask", such as the *Pudding Task*.[38] Say we have a countable set of disjoint regions $r_1, r_2, \ldots$ which together compose a finite region s, all in a world entirely free of pudding. At step one, fill r_1 with chocolate pudding; at step two fill r_2, and so on. When all of the steps have been carried out, s is entirely filled with chocolatey goodness. So the total amount of pudding in the world is $m(s)$. The total amount of pudding served is $m(r_1) + m(r_2) + \ldots$. So if $m(s)$ is greater than the sum, then there is more pudding in the world than was served. Whence this extra pudding?

Intuition presses us toward this principle: the sizes of the finite fusions should converge to the size of the infinite fusion. If we think of the regions involved as also forming a convergent sequence—the finite fusions converge to the infinite fusion—then there's another way to say this: intuitively, the measure function should be *continuous*.

Intuition favors continuous functions, yet continuity often fails in our dealings with the infinite.[39] Consider, for instance, the *Painting Task*. Suppose we have an infinite sequence of blank pages labeled with the natural numbers, and a large can of paint. At step one, paint the first page; at step two paint the second page, and so on. After any finite number of steps, there are infinitely many blank pages. Yet when the supertask is complete, there are no blank pages left at all. Again continuity fails: the limit of a sequence of infinities is certainly not zero, so the limit gives the wrong result for this supertask. Why shouldn't the Pudding Task be similar to the Painting Task in this respect? As the number of steps carried out approaches infinity, the amount of pudding approaches a particular amount—but when the number of steps reaches infinity, there is a different amount. There is a "break" at infinity. Surprising, but hardly impossible.

One might worry that completing such supertasks would make it possible to violate certain physical principles: a one-mile bridge could be built with only half a mile's worth of concrete, by building it up from the right set of parts. One available reply is that such physical principles *can* be violated: in worlds with the kind of space in question, it may be physically possible to build a bridge on the cheap in just this way. If the world has the right conservation laws, though, then the supertask cannot be carried out. It is a little puzzling that a law of physics might rule out an infinite sequence of operations without ruling out any finite subset of them, but this kind of thing happens. For instance, in a Newtonian world, conservation of momentum rules out an infinite sequence of collisions between vanishingly small balls at vanishingly small distances from one another, without ruling out any particular collision.[40] The physical law is inconsistent with the initial conditions that give rise to the infinite collision scenario—and so the law rules the scenario out. Similarly, in worlds where conserva-

[38] Forrest argues along these lines in Forrest, "How innocent is mereology?"; for a more elaborate example, see Hawthorne and Weatherson, "Chopping Up Gunk".

[39] This line of thought is influenced by conversation with Frank Arntzenius.

[40] Laraudogoitia, "A Beautiful Supertask".

tion of concrete is a law of nature, the bridge-building supertask can't be carried out. Whatever initial conditions might give rise to it are physically impossible.

One might also worry because measure *supervenience* fails. If the size of a region is not determined by the sizes of a collection of disjoint things that exhaustively compose it, doesn't that make its size a sort of spooky, free-floating thing? One reply is *tu quoque*: we are already accustomed to exactly this situation in pointy space, since an uncountable collection of point-sized regions may be rearranged to make up a region of any size whatsoever. If this is acceptable, we might equally well learn to live without size supervenience for countably infinite fusions of gunky regions. But the "spookiness" worry might be raised more generally: what fixes the size of a region, if not the sizes of its parts? For pointy space, the answer to this question is written in the annals of classical measure theory. For topologically gunky space it turns out to be a difficult question.

The natural place to look for a model for finitely additive, topologically gunky space is Tarski's open regular sets. Tarskian regions have the mereology and topology of their representative open sets; it would be nice to similarly fix the size of a region to be the (Lebesgue) measure of its representative. But this won't do.[41] To see why not, consider once more the Cantor Regions, labeled $a, b_1, b_2, c_1, c_2, c_3, c_4, \ldots$. Recall: the size of a is $\frac{1}{4}$, the combined sizes of the b's is $\frac{1}{8}$, the combined sizes of the c's is $\frac{1}{16}$, and so on; all of these together add up to $\frac{1}{2}$; and the fusion of the Cantor Regions is the unit interval. Now let the *Big Pieces* be a, the c's, the e's, etc., and let the *Little Pieces* be the b's, the d's, the f's, etc. Since the Big Pieces are twice as big as the Little Pieces, and all together they add up to $\frac{1}{2}$, the Big Pieces' sizes add up to $\frac{1}{3}$ and the Little Pieces' sizes add up to $\frac{1}{6}$. Now consider *Big Cantor*, the fusion of the Big Pieces, and *Little Cantor*, the fusion of the Little Pieces. How big are they? According to the proposed rule, the measure of Big Cantor is $\frac{1}{3}$ and the measure of Little Cantor is $\frac{1}{6}$.[42] But note: the fusion of Big Cantor and Little Cantor is the unit interval. So on this account of size, *Finite* Additivity breaks down!

Fixing sizes based on the open regular sets was a bad idea. The closed regular representatives do no better—the closed representative of Big Cantor has size $\frac{5}{6}$, and the closed representative of Little Cantor has size $\frac{2}{3}$; these don't add up to 1 either.[43] Appealing to a region's representative equivalence class is even less helpful: by adding

[41] Thanks to John Hawthorne for pushing me to get clear on this.

[42] Let I, $U_1, U_2, \ldots$, and $V_1, V_2, \ldots$ be the representative open sets for the unit interval, the Big Pieces, and the Little Pieces, respectively, and let $U^* = \bigcup U_i$ and $V^* = \bigcup V_i$. U^* is an open regular set. This is because every point in I is either (i) a member of U^*, (ii) a member of V^*, or else (iii) a boundary point of both U^* and V^*. So the closure of U^* is $I - V^*$, and the interior of $\mathrm{cl}\, U^*$ is U^*. Thus U^* is Big Cantor's representative. Furthermore, since Lebesgue measure is countably additive, $m(U^*) = \sum m(U_i) = \frac{1}{3}$. Similar considerations hold for Little Cantor.

[43] The closed representatives of Big Cantor and Little Cantor are the closures of their open representatives: namely $I - V^*$ and $I - U^*$, respectively.

and subtracting regions like the Cantor Dust, we can find members of each equivalence class with any given positive size. Which of these should we choose? The pointy representatives don't pin down the sizes of Big and Little Cantor.

Here's another way of characterizing this problem: Tarskian space includes regions that are not *Jordan-measurable*.[44] To define Jordan measure, we start with a collection of elementary regions whose sizes are already well understood—typically rectangles (in one dimension: intervals). We can calculate the measure of any finite fusion of disjoint rectangles by addition. Call these finite fusions the *blocks*. We then define the *outer measure* of a region x as the lower limit of the sizes of the blocks that contain x, and the *inner measure* as the upper limit of the sizes of the blocks that are part of x. If the outer and inner measure agree, then x is Jordan-measurable, and their common value is the Jordan measure of x. Jordan measure is finitely additive; in effect, Jordan measure starts with the sizes of elementary regions and unrolls the consequences of Finite Additivity for the rest. But there are regions whose sizes can't be determined in this way—and, sadly, Big Cantor and Little Cantor are among them. Big Cantor's inner measure is $\frac{1}{3}$, and its outer measure is $\frac{5}{6}$. In pointy space, Lebesgue measure captures these regions by going beyond finite collections of rectangles to *countable* collections—but since Countable Additivity fails in Tarski's space, this won't help.

One response to this failure is to cut the offending regions loose, either by admitting that Tarskian space includes peculiar regions that lack sizes, or else by restricting the space to include just the Jordan-measurable regions. The latter option involves giving up Unrestricted Composition for infinite collections of regions. These are costs we hoped to escape in gunky space, but at any rate they aren't much worse than the costs faced by points: these two options parallel the two responses to non-measurable sets in Section 3.1. For a model following the latter option, represent gunky regions with *Jordan-measurable open regular sets*, and assign regions the same parts, connection, and sizes as their representatives. The mereology of this space is not quite standard: while every *finite* set of regions has a fusion, some *infinite* sets do not: for instance, neither the Big Pieces nor the Little Pieces have a fusion.[46]

A second response to the failure of Jordan measure is to introduce "spooky" sizes which are not pinned down by elementary regions and additivity considerations. In fact, there are infinitely many finitely additive measures that extend Jordan measure to all of the Tarskian regions.[47] We can assign Big Cantor any size between its inner

[44]Jordan measure is to the Riemann integral as Lebesgue measure is to the Lebesgue integral.[45]

[46]The mereology forms a Boolean algebra, but not a *complete* Boolean algebra.

[47]This is an application of a general result by Tarski: if S is a subalgebra of a Boolean algebra A, then any finitely additive measure function defined on S has a finitely additive extension over all of A.[48]

If $a \in A$ and $a \notin S$, we can extend the measure on S to the larger subalgebra of elements of the form $(b \wedge a) \vee (c - a)$ (where $b \wedge a = b - (b - a)$). We set the size of such an element to be $\text{outer}_{m,S}(b \wedge a) + \text{inner}_{m,S}(c - a)$, where $\text{outer}_{m,S}(x) = \inf\{m(s) : x \leq s \in S\}$ and measure of $\frac{1}{3}$ and its outer measure of $\frac{5}{6}$, and go on to fix sizes for the rest of the unsized regions in a consistent fashion. All of the measure functions we obtain agree on simple cases like rectangles, and all of them obey additivity. What could favor one of these measure functions over any other? For instance, what might determine that Big Cantor is truly bigger than Little Cantor, rather than the other way around? The truth about the sizes of such regions might be a matter of "brute fact", but this introduces a host of brute facts where on the classical story about measure we have very few. There is, then, a model of Tarskian space with a finitely additive measure function; in fact, there are infinitely many non-isomorphic models. The problem is not that there is no way of assigning sizes—it is that there are too many.

This leaves us with two kinds of finitely additive space. The first kind requires giving up either Measurability or an infinite composition principle, and the second kind requires giving up a supervenience principle. These costs are considerable, but not unthinkable—remember, the point lover faces analogous compromises. Besides these costs, merely finitely additive sizes lead to counterintuitive discontinuities, but intuitions are untrustworthy when it comes to questions about continuity at infinity. Note finally: *uncountable additivity* (insofar as uncountable sums make sense) is a principle that pointy space has long since abandoned. Are *countably* infinite sets more sacred? Finitely additive space is closer to the Whiteheadian picture than either Arntzenius's or Forrest's alternatives—it is full-fledged gunk in the mereological, topological, and measure-theoretic senses. In it we have a reasonable candidate for at least some of the gunk theorist's purposes.

5 Conclusion

I have presented a strong argument against naïve Whiteheadian space: a small set of independently plausible conditions on gunky space turns out to be inconsistent. The gunk theorist can respond in several ways. The options on the table include Arntzenius's measure-theoretic gunk that violates the No Boundaries principle, Forrest's space that introduces revisionary measure theory by way of revisionary mereology, and a version of Tarski's geometry where sizes disobey Countable Additivity. Each of these options goes against certain intuitions: roughly, Arntzenius's space sacrifices an intuitive bit of topology, Forrest's space an intuitive bit of mereology, and Tarski's space an intuitive bit of measure theory. But I insist: the usual pointy models

$\text{inner}_{m,S}(x) = \sup\{m(s) : x \geq s \in S\}$. This is finitely additive and agrees with m on S. The measure is extended to the entire algebra A using an induction principle (When A is uncountable this relies on the Axiom of Choice).

Note that in this argument $\text{inner}_{m,S}(b \wedge a) + \text{outer}_{m,S}(c - a)$ yields a measure just as well, and any weighted average of measures is itself a measure. Therefore as long as there is some element of A whose outer and inner measures disagree, m has infinitely many finitely additive extensions.

of space violate intuition as well, as curiosities like the Banach-Tarski theorem (among many others) go to show. We have to live with a bit of counter-intuitiveness in any case.

How we choose among these theories will depend on what we are choosing them for. If the goal is a theory of space that is adequate for actual-world physics, Arntzenius's measure-theoretic gunk seems far and away the most satisfactory. If on the other hand the main concern is capturing common sense spatial reasoning, one of the finitely additive models may be better. And if the question is what spatial structures are broadly possible, then (short of further argument to the contrary) possibilities abound.

One lesson to draw from this discussion is a reminder that innocent-seeming premises can lead quickly into hidden incoherencies—we should heed Hume's admonition to "be modest in our pretensions".[49] An account may be intuitive without being consistent, especially when extrapolations to infinity are involved. But there are babies as well as bath-water in intuition's tub: while the naïve conception of gunk is problematic, the core idea displays a hardy resilience. In particular, my argument certainly does *not* rule out gunky space as a physically interesting candidate for our actual universe's space-time, as a metaphysical possibility, or as a consistent framework for spatial reasoning. Avoiding contradictions does force certain counterintuitive results on us—any theory of space has its puzzles and surprises—but we should expect that much.

References

Arntzenius, Frank. "Gunk, Topology, and Measure". In: *Oxford Studies in Metaphysics*. Ed. by Dean Zimmermann. Oxford University Press, 2008, pp. 225–247.

— "Is Quantum Mechanics Pointless?" In: *Philosophy of Science* 70.5 (2003), pp. 1447–1457.

Arntzenius, Frank and John Hawthorne. "Gunk and Continuous Variation". In: *Monist* 88.4 (2005), pp. 441–465.

Birkhoff, Garrett. *Lattice Theory*. Colloquium publications. Providence: American Mathematical Society, 1967.

Clarke, Bowman. "A Calculus of Individuals Based on 'Connection'". In: *Notre Dame Journal of Formal Logic* 22.3 (1981), pp. 204–218.

Forrest, Peter. "How innocent is mereology?" In: *Analysis* 56.3 (July 1996), pp. 127–131.

— "Mereological summation and the question of unique fusion". In: *Analysis* 56.3 (2007), pp. 127–131.

[49] Hume, *An Enquiry Concerning Human Understanding*, p. IV.ii.

Forrest, Peter. "Nonclassical Mereology and Its Application to Sets". In: *Notre Dame Journal of Formal Logic* 43.2 (2003), pp. 79–94.

Frink, Orrin. "Jordan Measure and Riemann Integration". In: *The Annals of Mathematics* 34.3 (1933), pp. 518–526.

Hawthorne, John and Brian Weatherson. "Chopping Up Gunk". In: *Monist* 87.3 (2004), pp. 339–350.

Hume, David. *An Enquiry Concerning Human Understanding*. Ed. by P. F Millican. Oxford: Oxford University Press, 2007.

Laraudogoitia, Jon Perez. "A Beautiful Supertask". In: *Mind* 105.417 (1996), pp. 81–83.

Lewis, David. *Parts of Classes*. Cambridge: Blackwell, 1991.

Munroe, M.E. *Introduction to Measure and Integration*. Reading: Addison-Wesley, 1953.

Nagel, Ernest. "Measurement: Mit deutscher Inhaltsangabe". In: *Erkenntnis* 2 (1931), pp. 313–333.

Röper, Peter. "Region-Based Topology". In: *Journal of Philosophical Logic* 26.3 (1997), pp. 251–309.

Sider, Ted. "Van Inwagen and the Possibility of Gunk". In: *Analysis* 53.4 (1993), pp. 285–289.

Sikorski, Roman. *Boolean Algebras*. Berlin: Springer Verlag, 1964.

Skyrms, Brian. "Logical Atoms and Combinatorial Possibility". In: *Journal of Philosophy* 90.5 (1993), pp. 219–232.

Tarski, Alfred. "Foundations of the Geometry of Solids". In: *Logic, Semantics, and Metamathematics*. Ed. by J.H. Woodger. Oxford: Clarendon Press, 1956, pp. 24–29.

Uzquiano, Gabriel. "Mereological Harmony". In: (Ms).

Varzi, Achille C. "Mereology". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Winter 2006. URL: http://plato.stanford.edu/archives/win2006/entries/mereology/.

— "Parts, wholes, and part-whole relations: The prospects of mereotopology". In: *Data and Knowledge Engineering* 20 (1996), pp. 259–286.

Zimmerman, Dean. "Could Extended Objects Be Made Out of Simple Parts? An Argument for 'Atomless Gunk'". In: *Philosophy and Phenomenological Research* 56.1 (1996), pp. 1–29.

— "Indivisible Parts and Extended Objects: Some Philosophical Episodes from Topology's Prehistory". In: *Monist* 79.1 (1996), pp. 148–180.