

The Fabric of Space: Intrinsic vs. Extrinsic Distance Relations

PHILLIP BRICKER

Start with ordinary Euclidean space. Be a realist about points, and about distances between points. I ask: how are the points interwoven to form the fabric of space? Are there direct ties only between "neighboring points," so that points at a distance are connected only indirectly through series of such direct ties? Or are there also direct ties between distant points, so that the fabric is reinforced, as it were, by irreducibly global spatial relations? To fix ideas, roughly, try the following thought experiment. Take a scissors and cut along some plane in space, severing the points just to one side of the plane from the points just to the other. Does space thereby fall into two disconnected pieces, so that points on one side now stand at no spatial distance from points on the other? Or does space, being reinforced, retain its shape?

If space is maximally reinforced by direct ties of distance, then a distance relation, such as *being twenty feet from*, is intrinsic to the points that stand in it. Whether or not the relation holds depends solely upon the intrinsic nature of the two points, and of the composite of the points. On the other hand, if space is not maximally reinforced by direct ties of distance, then a distance relation will not in general be intrinsic to the points that stand in it, and its holding may depend in part upon features of the surrounding space. What features? To fix ideas, roughly, try the following thought experiment. Start with two points, say, twenty feet apart, and remove some of the space directly between the two points. (I do not mean just the matter or energy occupying the space; I mean the space itself.) I ask: now how far apart are the points? Are they *still* twenty feet apart, on the grounds that distance, being intrinsic, is indifferent to changes in the intervening space? Or are they now *less* than twenty feet apart, on the grounds that there is now less space between them? Or are they now *more* than twenty feet apart, on the grounds that the shortest (continuous) path from one to the other is now more than twenty feet long? We have three competing answers, each, I think, with some intuitive appeal. Which is correct? And how can we tell?

The situation is familiar. We start with some notion from ordinary language or thought—in this case, the notion of distance (but compare the notions of person, cause, law, matter). We notice that there are different criteria associated with the notion, depending upon the context of application or of thought. But given the presuppositions under which the ordinary notion has evolved—in this case, presuppositions about the Euclidean nature of space—the different criteria fit together as well as you please. Then, driven perhaps by science, by mathematics, or by analytic philosophy, we consider extraordinary physical or logical possibilities that violate the presuppositions—for example, space with a “hole”—and the different criteria are seen to come apart. We are left with a plurality of competing conceptions, typically none of which captures all that was thought to be essential to the original ordinary notion. The question then arises: which conception should we accept?

It would be wrong in general to expect a univocal answer. Competing conceptions may be evaluated along at least three different dimensions. One can ask: which conception best corresponds to the ordinary notion with which we began? One can ask: which conception is mathematically, or philosophically, more fruitful, say, by leading to more interesting and powerful generalizations. Or one can ask: which conception has application at the actual world according to our best physical, or perhaps philosophical, theories?

In this paper, I evaluate various conceptions of distance. There are clear losers, but no clear winner, no conception that dominates the score on all dimensions of evaluation. I recommend pluralism: different conceptions can peacefully coexist as long as each holds sway over a distinct region of logical space. But when one asks which conception holds sway at the actual world, one conception stands out. It is the conception of distance embodied in differential geometry, the conception that underlies modern treatments of physical space (and spacetime) based upon Einstein’s general relativity. On this conception, all facts about distance are analyzed in terms of “local” facts about distances between “neighboring points.” Putting quantum mysteries to one side, I would say that this “local” conception gives the best account of distance at the actual world.¹ But there is a problem: the “local” conception, notwithstanding its mathematical and physical credentials, appears metaphysically suspicious. In the final section, I try to give the “local” conception a sound metaphysical footing.

A word of caution. My question whether distance relations are intrinsic to pairs of points should not be confused with the oft-discussed question whether “space has an intrinsic metric.” Reichenbach, Grünbaum, and many others held that facts about the congruence of intervals of space are imposed from the outside, as it were, by our conventions for interpreting the behavior of material rods or rays. Without these conventions, they held, there are no facts about congruence; with different conventions, there are different such facts. I simply reject this here. The question here is not whether to be a metrical realist or conventionalist, but rather, assuming realism, what sort of realist to be.²

I

I need some preliminary notions and assumptions. I will speak of possible worlds. In this paper, I restrict attention to worlds at which space exists, at which space is composed entirely of points, and at which there are determinate facts about the distance, say in feet, between pairs of points. For these worlds, both intraworld and transworld comparisons of distance are meaningful. For the most part, I also restrict attention to *Newtonian* worlds, worlds at which there are determinate facts about the “identity over time” of points of space, and at which distances between points do not change over time. Newtonian worlds need not have Euclidean space. I assume that a variety of spatial structures will be exhibited at Newtonian worlds: curved and flat, finite and infinite, with and without boundaries, continuous and discrete.³ Although I speak for simplicity primarily of spatial distance in Newtonian worlds, what I say applies more generally, *mutatis mutandis*, to temporal duration, and to intervals of spacetime in relativistic worlds. In section VI, the focus will switch from space to spacetime.

It is a matter of indifference whether one speaks of a distance *function* assigning non-negative reals to pairs of points, or of a multitude of distance *relations*, one for each non-negative real; I will speak of distance relations. I assume that the distance relations satisfy, at each Newtonian world, the usual constraints. Write ‘ $D_r(\mathbf{p}, \mathbf{q})$ ’ for ‘ \mathbf{p} is r feet from \mathbf{q} .’ Let ‘ r ’, ‘ s ’, and ‘ t ’ range over non-negative reals.⁴ Then, for any world, for any points \mathbf{p} , \mathbf{q} , and \mathbf{r} at the world:

- (D0) $D_r(\mathbf{p}, \mathbf{q})$, for exactly one r .
- (D1) $D_r(\mathbf{p}, \mathbf{q})$, iff $D_r(\mathbf{q}, \mathbf{p})$, for all r .
- (D2) $D_0(\mathbf{p}, \mathbf{q})$, iff $\mathbf{p} = \mathbf{q}$.
- (D3) If $D_r(\mathbf{p}, \mathbf{q})$, and $D_s(\mathbf{q}, \mathbf{r})$ and $D_t(\mathbf{p}, \mathbf{r})$, then $t \leq r + s$, for all r, s , and t .

I will freely apply mereology to points of space. Whenever there are some points, there is a unique *fusion* of those points. In particular, any pair of points, \mathbf{p} and \mathbf{q} , has a unique fusion, $\mathbf{p} + \mathbf{q}$. Whenever X is a proper part of Y , there is a unique *difference*, $Y - X$, which is the fusion of the parts of Y that do not overlap X . Space at a Newtonian world is the fusion of all the points of space existing at the world. I stay neutral as to whether a Newtonian world contains, in addition to the parts of space, entities that occupy those parts. If not, then the properties and relations ordinarily attributed to the “occupants” of parts of space must be attributed directly to the parts of space themselves. I assume that the spaces of distinct worlds do not overlap, that no point inhabits more than one world. Modal assertions that are *de re* points or regions of space must therefore be interpreted with respect to an appropriate counterpart relation.⁵

I assume that there are certain primitive or fundamental properties and relations, the holding or failing to hold of which suffices to determine, at any world, all the qualitative facts at that world. In particular, there are primitive or

fundamental spatial properties and relations which suffice to determine, at any world, all the facts about distance. Whether or not the distance relations are themselves among the primitives is a question soon to be addressed. I will call the primitive or fundamental properties and relations *perfectly natural* properties and relations.⁶ I assume that the part-whole relation is perfectly natural.

I will need to speak of worlds or parts of worlds being (intrinsic) *duplicates* of one another. I define 'duplicate' in terms of perfectly natural properties and relations:⁷ for all X and Y , X and Y are *duplicates* iff there is a one-one correspondence between the parts of X and the parts of Y that preserves all perfectly natural properties and relations. (Remember: everything is a part of itself.) I call any such correspondence establishing that X and Y are duplicates an (X, Y) -counterpart relation; to each part Z of X , it assigns a unique part W of Y to be its (X, Y) -counterpart. (I drop the prefix when context allows.) Note that, for any (X, Y) -counterpart relation, (X, Y) -counterparts are duplicates of one another. However, duplicate parts of X and Y will not be (X, Y) -counterparts, for any (X, Y) -counterpart relation, unless they are similarly related to the other parts of X and Y . Note also that, since there will in general be more than one (X, Y) -counterpart relation, a part Z of X and a part W of Y may be (X, Y) -counterparts relative to some (X, Y) -counterpart relations, but not others. Nonetheless, in presenting examples I will leave the relation unspecified, and say simply that Z and W are (X, Y) -counterparts. That will not lead to trouble because what I say will hold true for an arbitrarily chosen (X, Y) -counterpart relation, assuming, of course, that a single such relation is held fixed throughout the example.

I turn now to the notion of an intrinsic property or relation. Intuitively, a property is intrinsic just in case whether it holds of an object depends only upon the way the object is in itself. Let us take the way an object is in itself—its *intrinsic nature*—to be given by the disposition of perfectly natural properties and relations among the object and its parts. Then we have, in terms of duplicates: A property P is an *intrinsic nature* iff P is had by all and only the duplicates of X , for some X . And, since an intrinsic property of an object is one that depends only upon the object's intrinsic nature, we have: A property P is *intrinsic* iff, for all X and Y , if X and Y are duplicates, then X has P iff Y has P . Note that, on this notion of intrinsic, an object's *haecceity*—the property of being that object—is not one of the object's intrinsic properties, since it is not shared by the object's duplicates.

The notion of intrinsic can be extended to relations in a natural way. Let us say that a (dyadic) relation is *intrinsic* just in case whether or not it holds of a pair $\langle X_1, X_2 \rangle$ depends only upon the intrinsic natures of X_1 , X_2 , and the fusion $X_1 + X_2$. Then, we have in terms of duplicates: A (dyadic) relation R is intrinsic iff, for all X, X_1, X_2, Y, Y_1, Y_2 , if X and Y are duplicates, $X = X_1 + X_2$, $Y = Y_1 + Y_2$, X_1 a counterpart of Y_1 , and X_2 a counterpart of Y_2 , then R holds of $\langle X_1, X_2 \rangle$ iff R holds of $\langle Y_1, Y_2 \rangle$.⁸ (Similarly for relations of three or more places.) If a property or relation is not intrinsic, it is *extrinsic*. Note that it follows immediately from the definitions that the

perfectly natural properties and relations are themselves intrinsic. No surprise: the assumption was built into the definitions from the start.⁹

Finally, I will need two modest assumptions about the plenitude of possible worlds. One, a principle of recombination for points, I will introduce when I need it in section VII. The other is this: for any part of the space of a Newtonian world, there is a Newtonian world whose entire space is a duplicate of that part. Actually, I only apply this assumption to rather ordinary parts of a three-dimensional Euclidean space. To illustrate: start with an ordinary world satisfying the laws of Newtonian mechanics. The principle posits a world just like it except for a "hole" in space. At this world, conservation laws fail in the vicinity of the "hole". Objects entering the "hole" simply vanish; objects emerging from the "hole" appear out of nowhere. Bizarre, indeed. But logically impossible? Contemplate that as you drift towards the black hole at the center of the Milky Way!

II


I now present a multiple-choice exam. I invite the reader to try it. The questions involve precisely formulated variations on the thought experiments from the introduction. The original formulations were unsatisfactory. They were naturally understood to involve *de re* counterfactuals—e.g., if some of the space between two points were removed, those points would be closer together. The interpretation of such counterfactuals is not fixed once and for all: different contexts may favor different comparative similarity relations on worlds, and different counterpart relations on points.¹⁰ I had a particular interpretation in mind: only when so interpreted do responses to the thought experiments have the intended metaphysical consequences. Therefore, to rule out unintended interpretations, I bypass the counterfactual formulations of the thought experiments and speak directly in terms of possibilities. (Once the intended interpretation is well established, I will allow myself to slip back into the counterfactual mode.)

DISTANCE EXAM

Part I: Removing Space. Consider a world with a three-dimensional Euclidean space, E . Let p and q be points of E twenty feet apart. Let X_i (for i from 1 to 5) be a part of E that includes p and q . Consider a world whose entire space Y_i is a duplicate of X_i . Let p' and q' in Y_i be counterparts of p and q in X_i , respectively.

- (1) X_1 is $E - A$, where A is an open¹¹ infinite slab bounded by two parallel planes, ten feet wide, centered on and perpendicular to the line segment connecting p and q . (See figure 1.) How far apart are p' and q' in Y_1 ?


- (a) Twenty feet apart.
 (b) They stand in no distance relation.
 (c) Ten feet apart.

FIG. 1: CROSS-SECTION OF Y_1 .

- (2) X_2 is $E - B$, where B is an open sphere, ten feet in diameter, centered on the point midway between p and q . (See figure 2.) How far apart are p' and q' in Y_2 ?

- (a) Twenty feet apart.

- (b) $10\sqrt{3} + 5\pi/3$ feet apart.
 (Greater than twenty feet!)

FIG. 2: CROSS-SECTION OF Y_2 .


- (c) Ten feet apart.

- (3) X_3 is $E - C$, where C is an open right circular cone with height and radius each ten feet, with vertex on the point midway between p and q , and with axis perpendicular to the line segment connecting p and q . (See figure 3.) How far apart are p' and q' in Y_3 ?

- (a) Twenty feet apart.

- (b) Twenty feet apart.

- (c) $10\sqrt{2}$ feet apart.
 (Less than twenty feet!)

FIG. 3: CROSS-SECTION OF Y_3 .

- (4) X_4 is the surface of a sphere with diameter twenty feet and with p and q at opposite poles. How far apart are p' and q' in Y_4 ?

- (a) Twenty feet apart.

- (b) 10π feet apart.
 (Half the sphere's circumference.)


- (c) Zero feet apart.

- (5) X_5 is an infinite wavy plane whose hills and valleys are an alternating series of infinite half cylinders of diameter ten feet, with p and q on adjacent summits. (See figure 4.) How far apart are p' and q' in Y_5 ?

- (a) Twenty feet apart.

- (b) 10π feet apart.
 (A quarter of the cylinder's circumference, quadrupled.)

- (c) Zero feet apart.

FIG. 4: CROSS-SECTION OF Y_5 .

Part II: Adding Space. Consider a world with a two-dimensional space, X_6 , (2 or 7). Let p and q be points of X_6 , twenty feet apart. Consider a second world with a three-dimensional Euclidean space, E , and a part Y_6 of E that is a duplicate of X_6 . Let p' and q' in Y_6 be counterparts of p and q in X_6 respectively.

- (6) X_6 is the surface of a sphere with circumference eighty feet. How far apart are p and q in Y_6 ?

- (a) Twenty feet apart.

- (b) $40\sqrt{2}/\pi$ feet apart.
 (Length of chord subtending a quarter of a great circle.)

- (7) X_7 is a Euclidean plane. How far apart are p' and q' in Y_7 ?

- (a) Twenty feet apart.

- (b) Could be any distance d , $0 < d < 20$ feet, depending upon the nature of E and the choice of Y_7 .

III

Now for the answers. Unfortunately, no one answer key will do. Different conceptions of distance answer the questions differently. Consider first the conception according to which distance relations are intrinsic to the pairs of points that stand in them. Call this the *intrinsic* conception of distance. (Presumably—though nothing rests on it—distance relations are not only intrinsic on this conception, but perfectly natural; for what other intrinsic features of the two points or their fusion could suffice to determine the distance between them?) On the intrinsic conception of distance, if points p and q are twenty feet apart, and p' , q' , and $p'+q'$ are duplicates of p , q , and $p+q$, respectively, then p' and q' are twenty feet apart. Now, for all seven questions, p' , q' and $p'+q'$ are counterparts of p , q , and $p+q$, respectively; and counterparts are duplicates; so, on the intrinsic conception, the answer seven times over is: (a) twenty feet; "Additions to" and "removals from" the space surrounding two points are nowise relevant to the distance between them.¹²

Central to the intrinsic conception is the notion of *congruence*, generalized to apply to parts of space perhaps from different worlds: X is *congruent* to Y iff there is a one-one correspondence between the points of X and the points of Y that preserves all distance relations. On the intrinsic conception, duplicate parts of space are congruent. I assume the intrinsic conception accepts a partial converse as well: congruent parts of space are *spatial* duplicates, that is, they agree with respect to all their intrinsic *spatial* properties.¹³ Thus, on the intrinsic conception, congruence serves to delimit the border between the intrinsic and the extrinsic.

The mathematical embodiment of the intrinsic conception is the abstract structure of a metric space. A metric space consists of a non-empty universe of points together with a family of distance relations (or a single distance function—it matters not) satisfying the axioms for distance listed above. The distance relations are taken as primitive, and other features of space—e.g., topological features—are defined in terms of the distance relations. The notion of a metric space is mathematically simple, yet extremely general: it encompasses spaces that are curved, flat, continuous, discrete, and all manner of hybrids thereof.¹⁴ Since the distance axioms all quantify only universally over points, any part of a metric space, and so any duplicate of that part, is itself a metric space. That ensures that one may speak without impropriety of distances between points in the spaces Y_1 through Y_7 .

Now consider a second conception of distance: the distance between two points of a space is given by the length of the (or a) shortest continuous path through space from one of the points to the other. (More exactly: the greatest lower bound of the lengths of continuous paths from one to the other, since in general there need be no least length, but I henceforth ignore this complication.) Call this the *Gaussian* conception of distance. The paths through space are themselves parts of space, fusions of points.¹⁵ On the Gaussian conception, the assignment of lengths to paths is prior to the assignment of distances to pairs of points. The Gaussian characterization of distance may have an air of circularity about it; but the air is only apparent. One metrical notion—distance between points—is defined in terms of another metrical notion—length of paths. There is no attempt to analyze away all metrical notions. Later we shall ask how length of path can itself be analyzed; and then, of course, we shall have to be careful not to close the circle.

On the Gaussian conception, I suppose, the length of a path through space is an intrinsic property of that path. However, the distance between two points turns out not to be an intrinsic relation of those points. If some of the space surrounding two points is “removed,” some or all of the paths connecting those points may no longer exist, and the length of the shortest remaining path—the new distance between the points—may be greater than it was, or not defined. If the space surrounding two points is embedded in a larger space, new paths connecting the points may come into existence, and the length of the shortest connecting path—the new distance between the points—may be less than it

was. In short: the distance between two points does not depend solely upon the intrinsic nature of the fusion of the two points.

To illustrate, turn to the exam. The Gaussian conception answers (b) seven times over. In question (1), none of the paths connecting p and q in E have counterparts in Y_1 . Thus no path in Y_1 connects p' with q' , and the distance between them is undefined (equivalently, ∞ , given our convention). Y_1 is composed of two “island universes” with p' and q' inhabiting different islands.¹⁶ In question (2), the straight-line path from p to q in E has no counterpart in Y_2 . The shortest path from p' to q' in Y_2 is one that follows a tangent from p' to the edge of the hole, hugs the hole for a sixth of a turn, and then follows a tangent back to q' . (See figure 2.) Hence, the answer: $10\sqrt{3} + 5\pi/3$ feet apart. In question (3), the straight-line path from p to q in E has a counterpart in Y_3 , and the distance between p' and q' in Y_3 is the same as the distance between p and q in E : twenty feet. In question (4), the shortest path connecting p' and q' in Y_4 is half a great circle. In question (5), the shortest path connecting p' then up another quarter circle hill, around a half circle valley, and path between p' and q' is a counterpart, not of the straight-line shortest path between p and q in E , but of a longer, more circuitous path between p and q whose length is given by the answer (b). I consider questions (6) and (7) below.

Central to the Gaussian conception is the notion of an *isometry* between parts of space: X and Y are *isometric* iff there is a one-one correspondence between the points of X and the points of Y that (when extended to fusions of points) preserves lengths of paths. (More exactly, the image of a path with endpoints p and q is a path of the same length with endpoints the image of p and the image of q .) Since Gaussian distance is defined in terms of lengths of paths, isometries preserve Gaussian distance as well. Duplicate parts of space are isometric, since the length of a path is intrinsic. With that the intrinsic conception can agree. But the Gaussian conception accepts, whereas are spatial duplicates, and thus agree with respect to all their intrinsic spatial properties. On the Gaussian conception, isometries delimit the border between the intrinsic and the extrinsic: spatial properties are intrinsic just in case they are preserved by isometries, just in case they are *isometric invariants*.

To illustrate the difference between congruence and isometry, consider a “flat” plane F and a “way” plane W (such as X_5), each embedded in a three-dimensional Euclidean space E . F and W are not congruent, since no one-one correspondence between F and W can preserve the distances among four points of W not co-planar in E . Therefore, on the intrinsic conception, F and W are not spatial duplicates. However, F and W are isometric. Intuitively, this is because something flat, such as a piece of paper, can be made to wave without stretching or tearing, and thus without changing the lengths of any paths confined to its surface. Thus, on the Gaussian conception, F and W are spatial duplicates, and agree with respect to all their intrinsic spatial properties.

Let us now see how the Gaussian conception answers questions (6) and (7). For question (6), the Gaussian reasons: X_6 and Y_6 are duplicates; therefore isometric. The only parts of E isometric to X_6 are themselves surfaces of spheres with circumference eighty feet.¹⁷ Therefore Y_6 is one such. Since the shortest path between two points on the surface of a sphere is part of a great circle, and p and q are twenty feet apart, p and q are connected by a quarter great circle in X_6 . By the isometry, p' and q' are connected by a quarter great circle in Y_6 . The distance, then, between p' and q' in Y_6 , and so in E , is the length of a "wormhole" through the interior of the sphere; namely, the length of the chord that subtends the quarter great circle connecting p' and q' . This length, which is less than twenty feet, is given by answer (b): $40\sqrt{2}/\pi$ feet apart.¹⁸ For question (7), the Gaussian reasons: X_7 and Y_7 are duplicates; therefore isometric. Since E contains both flat planes and wavy planes that are isometric to X_7 , Y_7 may be either wavy or flat. If Y_7 is a wavy plane, then p' and q' are closer together in Y_7 (and E) than p and q are in X_7 : p' and q' are connected by a (straight-line) "wormhole" in E . How close together? The distance between p' and q' is just the length of the (straight-line) "wormhole," which may have any value in feet greater than zero and less than twenty, depending upon the "wavelength" of Y_7 . Thus, the Gaussian answers (b) to question (7).

The Gaussian conception of distance finds mathematical expression in the development of differential geometry. Here is how the conception is typically motivated.¹⁹ Start with a two-dimensional surface X embedded in a three-dimensional Euclidean space E . Ask: what geometrical features of the surface X could be ascertained by two-dimensional geometers whose measurements were entirely confined to X ? These features comprise X 's "intrinsic geometry."²⁰ The length of a path confined to X can be ascertained to any specified degree of accuracy by placing sufficiently small measuring rods end to end along the path; so lengths of paths are intrinsic to X . Moreover, features that depend only upon lengths of paths—the isometric invariants—could all be ascertained by the geometers, and so are intrinsic to X ; this includes, most famously, the Gaussian curvature at a point. But the true Euclidean distances between points of X are not intrinsic to X : they cannot be ascertained without "leaving the surface." The geometers do, however, have an intrinsic substitute for distance between points: *distance-within- X* , that is, the length of a shortest path in X between the points. Indeed, if the geometers (wrongly) take X to be all of space, they will (wrongly!) take *distance-within- X* to be the true distance.²⁰ The next step is to do away with the embedding space E , to consider a surface Y intrinsically just like X , but not embedded in any larger space. For this surface Y , the intrinsic geometry is all there is to geometry, and *distance-within- Y* is all there is to distance. The final step generalizes the Gaussian conception to apply to surfaces not embeddable within Euclidean space, and to spaces of dimension greater than two. In particular, since Euclidean distances are identical with *distances-within- E* , the Gaussian conception applies to E as well.

I will attempt to evaluate the motivation that underlies the Gaussian conception in section V. First I want to consider a third conception of distance. It is not, in my view, a serious contender. But I dare not ignore it: the most popular answer by far to question (1) is neither (a) nor (b), but (c). Call it the *naïve* conception. The leading idea is this: the distance between two points should be a measure of the amount of space between the points; but, unlike the Gaussian conception, the amount of space between two points need not be identified with the length of any continuous path. Indeed, although the naïve conception agrees with the Gaussian conception that distance relations are not intrinsic, the dependence of distance on the surrounding space is reversed. On the naïve conception, if space is "removed" from between two points, the points will then be *closer together*; if space is "added" between two points, the points will then be *farther apart*. How might the amount of space between two points be measured so as to capture these naïve intuitions?

Version 1. Many who answer (c) to question (1) have in mind closing up the gap left by the missing slab, and "stitching" the two remaining parts of Y_1 back together. The amount of space between p' and q' in Y_1 is then the length of the straight-line path in the stitched-up space. But that won't do. For one thing, on topological grounds the stitching cannot be seamless. One cannot avoid the seam by identifying boundary points on opposite sides of the gap; for that would violate the supposition that X_1 and Y_1 are duplicates (assuming the topological property *being connected* is intrinsic, and so must be shared by duplicates). But a seam composed of distinct, co-present boundary points would violate the distance axiom (D2) requiring that distinct points be some positive distance apart. How bad is that? *Perhaps* violations of (D2), if restricted to boundary points, could be tolerated on the grounds that boundary points may in effect be contiguous, and thus no distance apart.²¹ In any case, the stitching idea does not generalize. Depending upon the shape of the part of space removed, the stitching might be done in any number of ways, the choice among which is arbitrary; indeed, this is so for Y_2 through Y_5 . Thus, version 1 cannot in general provide determinate answers to questions about distance. We need another idea.

There are, however, ways to make the naïve conception more precise. I will consider the two most promising. For each, I assume that the length of a path is an intrinsic property of that path; and I will speak of "disconnected paths," whose lengths (when defined) are the sum of the lengths of their connected parts. *Version 2.* Determine the distance between p' and q' in Y_i (i from 1 to 5) as follows: start with the straight-line path in E connecting p and q ; take the part of this path, perhaps disconnected, that overlaps X_i ; then take as answer the length of this perhaps disconnected path (equivalently, of its counterpart in Y_i). When applied to questions (1) and (2), the result is answer (c): ten feet apart. But consider question (3). On version 2, the answer would be twenty feet, in agreement with the intrinsic and Gaussian conceptions. But someone who holds the naïve conception could say that p' and q' are closer together in Y_3 than p and q are in E . They could take the amount of space

between p' and q' to be the amount of space between p' and the cone-shaped hole plus the amount of space between the cone-shaped hole and q' . (See figure 3.) This illustrates version 3. Determine the distance between p' and q' in Y_i as follows: start with all the (continuous) paths in E' connecting p and q ; for each such path, take the perhaps disconnected part of the path that overlaps X_i ; then take as answer the least of the lengths of these perhaps disconnected paths (or, equivalently, of their counterparts in Y_i). (More generally, take the greatest lower bound of the lengths.) When applied to question (3), version 3 gives the answer (c): $10\sqrt{2}$ feet apart.

In the cases considered, either version 2 or 3 may appear to give plausible answers. But trouble looms for both. First, note that on either version distinct boundary points of a space may be zero feet apart in violation of axiom (D2). For example, two points on opposite "shores" in Y_1 , or two points on the edge of the "hole" in Y_2 . Perhaps, as noted above, that is tolerable. But now look more closely at how version 2 applies to Y_2 : Since boundary points r' and s' in Y_2 are zero feet apart (see figure 2), we have $D_5(p', r')$, $D_0(r', s')$, and $D_5\sqrt{3}(p', s')$ in violation of the triangle inequality (D3). Not so tolerable. No family of relations that violates the triangle inequality deserves to be called a family of distance relations.

That leaves version 3. Version 3 satisfies the triangle inequality, in much the same way as the Gaussian conception, by defining distance as the length of a shortest "path" (under an expanded notion of path). Since the shortest "path" in Y_2 from p' to s' goes by way of r' , $D_5(p', s')$ on version 3. But when applied to Y_4 and Y_5 , version 3, and version 2 to boot, fail to give plausible answers. Let's focus upon Y_4 , the duplicate of the surface of a sphere. On both versions 2 and 3, not only p' and q' , but any two points of Y_4 are assigned the distance: zero feet apart. That violates axiom (D2) in a big way. It obliterates all distinctions of distance, treating Y_4 in effect as a "space" with but a single point. Moreover, Y_4 is a space without boundaries: no path in Y_4 abruptly comes to an end. Even were one to tolerate violations of (D2) for boundary points, I see no comparable grounds for leniency here. I conclude that neither version 2 nor version 3 gives acceptable answers to questions about distance between points of Y_4 .

Could some fourth version of the naive conception answer any differently? The distance between p' and q' in Y_4 must be less than or equal to twenty feet, on the naive conception, because space was "removed" from E' ; yet the distance must be a measure of the amount of space between p' and q' . How could any such distance, other than zero, be singled out over any other? I reject the naive conception.

IV

That leaves two contenders: the intrinsic and the Gaussian conception. In this section, I question the Gaussian demarcation between intrinsic and extrinsic spatial properties. First, I ask the Gaussian whether the shape of a thing

is intrinsic to that thing. Suppose it is. Consider a wavy plane and a flat plane embedded in a three-dimensional Euclidean space. They are isometric; therefore they are spatial duplicates, says the Gaussian; therefore, since shapes are intrinsic, they have the same shape. That is plainly wrong: one is curved and the other is flat! So the Gaussian must deny that shapes are intrinsic. That looks bad. We ordinarily take the shape properties to be the very paradigm of intrinsic properties, of properties that depend only upon the way something is in itself. And we ordinarily would say that two things cannot be duplicates of one another unless they have the same shape. It appears that the Gaussian conception clashes with our ordinary ways of thinking about shape, and so, derivatively, about distance. The intrinsic conception, on the other hand, can uphold the intuition that shapes are intrinsic, since the wavy plane and the flat plane are not congruent with one another.

I suppose a Gaussian might respond as follows. Ordinary intuition, rightly understood, does not conflict with the Gaussian conception. Our intuitions about shape apply only to the objects of our experience, not to mathematical abstractions therefrom; and the objects of our experience are three-dimensional and Euclidean. Now, for ordinary three-dimensional parts of a three-dimensional Euclidean space—spheres, cubes, even paper-thin sheets—the parts are isometric only if they have the same shape. That allows the Gaussian to hold that shapes are "intrinsic" in a restricted sense: for ordinary three-dimensional parts of a three-dimensional Euclidean space, duplicate parts always agree in intuition. And ordinary intuition demands no more. Just as the domain of ordinary intuition is restricted to objects of experience, so should the sense in which ordinary intuition takes shapes to be intrinsic similarly be restricted.

This response lacks conviction. The two-dimensional surfaces of three-dimensional things, no matter how "abstract" or "ideal," seem to be objects of our intuition no less than the three-dimensional things themselves; and intuition pronounces the shapes of the former intrinsic no less than the latter. Nor does it much help to note that many two-dimensional surfaces, such as that of a sphere or a cube, are rigid, and so are isometric only if they have the same shape; for rigidity plays no role in the relevant intuition. The Gaussian should concede the clash with ordinary intuition. If no more condemns the Gaussian conception than, say, the acceptance of continuous paths through space that are nowhere "smooth"—once thought monstrous by intuition—condemns the standard mathematical analysis of continuity. Ordinary intuitions about matters susceptible to mathematical precision have often been found in need of revision. That is a small price to pay for the power and generality conferred by mathematics, or for the explanatory and predictive success of scientific theories mathematically based. Adherence to ordinary intuition is to some extent necessary to keep our bearings; but when mathematics gives us clear vision above and beyond, we should not hesitate to change our course. Agreement with ordinary intuition, by itself, favors the intrinsic conception but little.

V

In this section, the Gaussian takes the offensive. Suppose the geometers of a world have access to every part of their space. They have measured the length of every path, the area of every surface, the volume of every region. Then, says the Gaussian, they have the wherewithal to know all there is to know about the structure of their space: there are no spatial facts that are in principle inaccessible to geometric measurement. The intrinsic conception, we have seen, must disagree. Consider again a world with a two-dimensional space isometric to the Euclidean plane. On the intrinsic conception, the space may be "wavy" or "flat," depending upon facts purportedly about "distance"; but these "distance" facts are inaccessible to geometric measurement, even assuming the geometers have access to every part of their space.

Now the Gaussian objects: these facts are mysterious. If the two-dimensional space were embedded in some inaccessible higher-dimensional space, the "distance" facts could be understood in terms of inaccessible facts about the nature of the embedding. But by assumption there is no such embedding space. Rather, the inaccessible "distance" facts reflect the possibilities of embedding, what would have been the case, had the space been embedded in some higher-dimensional space. But to explain such facts in terms of the possibilities of embedding is to reverse the true order of things. What reason could there be, then, for taking these facts which are inaccessible to geometric measurement to be facts about *distance*, to be *spatial* facts? By maintaining that these inaccessible facts are facts about distance, the intrinsic conception posits a *phantom embedding space*, the ghost of a departed embedding.

Note that this argument does not rest upon a positivist premise. The Gaussian need not deny that there could be perfectly natural relations between points of space—even relations satisfying the distance axioms—knowledge of which is inaccessible to geometers who have access to all parts of space. The Gaussian need only deny that such relations could be the *distance* relations.

Before attempting to evaluate the argument, I want to examine the assumption that lies at its heart. It is a supervenience thesis. When stripped of colorful talk of tiny geometers, it comes to this: if the spaces of two worlds are isometric, then the spaces are congruent as well. In short: distances supervene upon lengths of paths.²² This assumption needs to be qualified in at least two ways. The alleged supervenience is contingent, not logical.

First, consider worlds with discrete space. In a discrete space, there are no continuous paths between points, so no lengths of continuous paths. Therefore, any two discrete spaces with the same number of points agree vacuously on all lengths of continuous paths. But the spaces need not agree on all distances between points, say, by having each point be an island universe all to itself. That is one possibility, I suppose. But I also suppose the points of a discrete space may stand in various distance relations, as long as the relations satisfy the axioms for distance. Indeed, I suppose there could be physical evidence that actual space (or spacetime) is discrete, and thus that the actual distance relations

are not Gaussian. Since the Gaussian analysis of distance cannot account for the variety of possible discrete—more generally, disconnected—spaces, the Gaussian assumption must be qualified: distances supervene upon lengths of paths, *for worlds with continuous space*.

Gaussian supervenience also fails, I think, at some worlds with continuous space. Consider worlds with "action at a distance": worlds at which forces act directly from one point to another without being propagated along a continuous path connecting the points. At such worlds, distance relations need not be Gaussian. For example, consider again Y_2 , the space with a "hole." Suppose that have been well confirmed: the force of gravity produces an acceleration that varies inversely with the square of the distance. Then, a measurement of the acceleration of objects located at p' and q' could give evidence against the Gaussian conception. Moreover, the "action at a distance" need not be instantaneous. Suppose that away from the "hole" it is well confirmed that all causal signals travel no greater than the speed of light. Then, a measurement of the time it takes signals to travel from p' to q' could give evidence against the Gaussian conception. The most the Gaussian is entitled to claim is this: there could be no evidence against the Gaussian conception at worlds where it has been established that all action is local, that is, propagated locally along continuous paths. For only in such worlds must evidence for distance relations be, *ipso facto*, evidence for lengths of paths.

The Gaussian supervenience thesis thus applies at most to *local-action* worlds, to worlds with continuous space and no action at a distance. The Gaussian argument against the intrinsic conception must similarly be limited in scope. (The argument erred, in particular, by focusing too narrowly upon what would be accessible to *geometers*, rather than what would be accessible to physicists more generally.) It follows that the intrinsic conception of distance cannot be jettisoned from logical space. The Gaussian must accede perfectly natural: at some worlds, distance relations are intrinsic, and presumably to the Gaussian analysis. Under pluralism, worlds with Euclidean space may have either intrinsic or extrinsic distance relations. That points up a flaw in my distance exam. The embedding space E was underspecified. Interpreted one way, the answer is (a) throughout; interpreted the other way the answer is (b). I shall have to give a lot of 'A's.

Pluralism is the best the Gaussian can hope for. Unfortunately, the Gaussian is not yet in a position to demand his share of logical space. Limiting the scope of the Gaussian argument undermines its force altogether. Any world at which Gaussian supervenience fails—be it a world with discrete space or less than a world whose space is a "wavy plane." The argument began as a general indictment of worlds with a phantom embedding space, no general indictment of worlds with a phantom embedding space. What remains is a specific indictment—based none too clearly on considerations of physical evidence—of *local-action* worlds with a phantom embedding space. But the

intrinsic conception is not committed to such worlds. If one takes a local-action world with intrinsic distance relations, and one "removes," say, a sphere from its middle, one gets a world with a phantom embedding space all right, but the world is no longer a local-action world. I conclude that the Gaussian is left without any argument against the intrinsic conception. For the sake of uniformity, why not take the intrinsic conception to apply to all worlds with distance relations?

VI

Here's why. Our best physical theory of space and time, Einstein's general relativity, is based upon differential geometry, and is Gaussian through and through. I suppose general relativity is logically possible, that is, true at some possible worlds. At these worlds, distance relations are Gaussian, not intrinsic. Moreover, to whatever extent we believe that general relativity is true at the actual world, to that extent we should believe that actual distance relations are Gaussian.

Before turning to the treatment of distance in general relativity, we need to further develop the Gaussian conception. Thus far, length of path has been left unanalyzed. If length of path is taken as primitive, then the Gaussian and the intrinsic conception are both *global* conceptions of distance: both apply primitive metrical notions to pluralities of points, in one case to paths, in the other to pairs. I now want to develop a *local* version of the Gaussian conception according to which the only primitive metrical notions are local properties of points. (I postpone an exact definition of 'local property' until the next section.) Let us again assume that the points of space have a *manifold* structure in terms of which paths can be characterized as continuous and smooth. Now, assign to each point of space a metric tensor \mathbf{g} , or $d\mathbf{s}^2$, which supplies information about distances within an "infinitesimal neighborhood" of the point; call the tensor \mathbf{g} at a point \mathbf{p} the *local metric* at \mathbf{p} .²³ Given two points \mathbf{p} and \mathbf{q} (no matter how "close together"), the distance between them is not determined by the local metric at \mathbf{p} and the local metric at \mathbf{q} . But given a path from \mathbf{p} to \mathbf{q} , the length of that path is determined by the local metric at each point along the path: it is the result of *integrating* $d\mathbf{s}$ along the path from \mathbf{p} to \mathbf{q} .²⁴ (In effect, the local metric at a point provides a set of "infinitesimal measuring rods," one for each direction, to be used for determining lengths of infinitesimal portions of paths passing through the point; integration then corresponds to measuring the length of a path by laying (continuum-many!) appropriately directed measuring rods end to end.) On the local Gaussian conception, the properties of having such-and-such local metric are taken as primitive. These properties then suffice to determine the length of any path through space, and so the Gaussian distance between any two points. There are no primitive global metrical properties or relations.

Now, I claim that general relativity is a local Gaussian theory. (Of course, here the local metric has Lorentz signature, since it provides information about infinitesimal intervals in spacetime, rather than infinitesimal distances in space;

the spacetime interval between two points is given by a longest, rather than a shortest, path.) A key insight behind general relativity is that all physics takes place by local action.²⁵ When applied to gravitation, this leads to Einstein's field equation, the fundamental law of gravitation in general relativity. Einstein's equation— $\mathbf{G} = 8\pi\mathbf{T}$ —tells how the local mass-energy density, given by the stress-energy tensor \mathbf{T} , relates to the local curvature of spacetime, given by the Einstein curvature tensor \mathbf{G} . (\mathbf{G} is analyzable in terms of the metric tensor \mathbf{g} .) More generally: when the fundamental laws of physics are formulated within general relativity, the metrical notions that occur in the laws are all *local* metrical properties, including the metric tensor \mathbf{g} . (This contrasts sharply with Newtonian physics: according to the fundamental law of gravitation, the force of gravity varies inversely as the square of the *distance*.) Now, I suppose that all and only the perfectly natural properties and relations instantiated at a law-governed world occur in the fundamental laws of that world. It follows that the properties of having such-and-such local metric are perfectly natural properties instantiated at general relativistic worlds, and that general relativity, as formulated by Einstein, is a local Gaussian theory.

I do not deny that one could give an empirically equivalent reformulation of general relativity in terms of global metrical relations; under specifiable conditions, global and local metrical relations are *interdefinable* with aid of the calculus. But that would not show global metrical relations to be perfectly natural, any more than, say, a reformulation of the laws of color (if there were such laws) in terms of grue and bleen would show grue and bleen to be perfectly natural. In general relativity, the local metric at a point is a *dynamic object*, a *prime mover*: it tells objects at that point how to move. It has the same claim to perfect naturalness as the other prime movers, such as the electromagnetic field at a point. It would be arbitrary and absurd to hold that some prime movers are perfectly natural, but not others.

It is well known that Einstein's general relativity eliminates primitive action at a distance. I have been arguing what is perhaps less well known, that general relativity eliminates primitive "distance at a distance." The reduction of global relations to local properties in general relativity applies to metrical relations as well. This has implications for the formulation of philosophically interesting supervenience theses. David Lewis has defended the viability of *Humean supervenience*, according to which all facts, other than facts about spatiotemporal distance, supervene upon local matters of particular fact. At worlds of Humean supervenience: "We have geometry: a system of external relations of spatiotemporal distance between points. . . . And at the points we have local qualities: perfectly natural intrinsic properties which need nothing bigger than a point at which to be instantiated. . . . All else supervenes on that."²⁶ But have we not just seen that, at least at local Gaussian worlds, even the relations of distance between points supervene on local matters of fact? Do we have, then, a sweeping elimination of all primitive global notions at local Gaussian worlds, a grand supervenience of the global on the local?

That would be too much to ask. On the local Gaussian conception, global distance relations supervene not on local metric alone, but on local metric plus manifold structure. Without manifold structure, no integration. Without integration, no analysis of global distance relations in terms of local metric. Manifold structure is in part topological structure, and topological structure, it is easy to see, is irreducibly global. Consider a two-dimensional Euclidean plane and (the surface of) an infinite cylinder. They are locally indistinguishable: each consists of continuum-many points that are locally Euclidean. But the plane and the cylinder differ topologically. For example, the plane, but not the cylinder, is *simply connected*: all closed paths can be continuously contracted to a point. Just as there are irreducibly global topological features of *space*, so also of *spacetime* at relativistic worlds.

Thus, general relativity suggests no grand supervenience of everything on local matters of particular fact; it suggests something more modest. Call it *Einsteinian supervenience*. At worlds of Einsteinian supervenience: we have a manifold of spacetime points (with topological and differential structure), and a distribution of perfectly natural local properties (including local metrical properties) over those points; all else supervenes on that. Of course, Einsteinian supervenience, like its Humean cousin, is philosophically controversial. I here claim only that it is the right supervenience thesis to consider at general relativistic worlds. Under pluralism, Einsteinian and Humean supervenience are not in conflict. Each holds contingently, and governs its own region of logical space. These regions differ with respect to the instantiation of perfectly natural spatiotemporal relations of distance. Given the success of general relativity, I suspect we are nearer to, if not within, the region of Einsteinian supervenience.

VII

I have argued that modern physics of spacetime is based upon a local Gaussian conception of (spatiotemporal) distance. In this final section, I ask whether the local Gaussian conception is metaphysically suspect. I claim that the local metric at a point, as standardly conceived, is not an intrinsic property of the point.²⁷ Thus, the local Gaussian appears to be committed to perfectly natural, *extrinsic* properties. That would introduce necessary connections between distinct co-inhabitants of local Gaussian worlds, namely, between points and their surrounding space; it would violate a modal "principle of recombination" that I, for one, would be loath to give up. I take this as a challenge not to the physicist—I am not so bold—but to the metaphysician: provide a coherent metaphysical foundation for modern space-time theories.

First, we need a precise characterization of local properties. The characterization requires topological structure. Say that a part of space, N , is a *neighborhood* of a point p iff some part of N includes p and is open in the topology of the space. (For example, in a three-dimensional Euclidean space, N is a neighborhood of p iff N includes some *open ball* around p , that is, all the points less than some positive distance r from p .) A property of points

P is local iff, for any points p and q , for any neighborhood N of p and any neighborhood M of q , if N is a duplicate of M and p is a (N , M)-counterpart of q , then P holds of p iff P holds of q . Note that if a property of points is intrinsic, then it is local; for counterparts, being duplicates, share all their intrinsic properties. But in general local properties need not be intrinsic. Call a property of points that is local but not intrinsic *neighborhood-dependent*.²⁸

The most familiar examples of neighborhood-dependent properties come from elementary calculus: derivatives of functions. Consider the position of some point-sized object as a function of time; suppose at time t it is located at point p . The *instantaneous velocity* of the object at t is the derivative of the position function evaluated at t . This derivative at t depends upon the object's position not only at t , but also at "neighboring" times. Or, turning this around, the derivative at t depends upon when the object is located not only at p , but also at "neighboring" points. The object's instantaneous velocity at t is thus a neighborhood-dependent property of both the time t and the point p . (In by a vector, and so instantaneous velocity is a vector as well, having both a magnitude (speed) and a direction (if non-zero). Both an object's speed and its direction of motion are neighborhood-dependent properties of points and of times.

Now, I claim that the local metric at a point p , as characterized in differential geometry, is a neighborhood-dependent property of p . That is because the local metric at p is an inner product on the tangent space at p : it takes a pair of tangent vectors as input, and gives a real number as output. (If the inputs are one and the same, the output is the squared length of the tangent vector.) The tangent vectors at p are defined as the *derivatives* of "smoothly" parametrized paths through p . (A parametrized path is a function from an interval of real numbers to the points of the path. If one thinks of the parameter as "time," then a "smoothly" parametrized path through p is a trip through p with no jolts or stops, and the tangent vectors at p are all the possible "velocities," or "states of motion," when passing through p .) These tangent vectors, being derivatives, give information not just about p , but about the space immediately surrounding p . For example, the dimensionality of the tangent space is the dimensionality, not of p which is zero, but of the immediately surrounding space. In short: the tangent vectors provide neighborhood-dependent information about p . Since the local metric at p is an operator on tangent vectors, it inherits neighborhood-dependence from its operands.²⁹

Thus, the local metric at a point, as standardly conceived in differential geometry, is neighborhood-dependent; and that is trouble for the local Gaussian conception. For, on the local Gaussian conception, the local metric is also perfectly natural. Apply the definitions from section I. If perfectly natural, neighborhood-dependent and intrinsic. *Contradiction*.

Perhaps we should revise our definitions, not our conception of distance. It was simply built into the definitions that all perfectly natural properties

are intrinsic. What was built in can be built out. For a cost. We need to introduce a primitive distinction between the perfectly natural properties that are intrinsic and those that are not. (So we no longer can analyze 'intrinsic' just in terms of perfectly natural properties and relations.) Now the definition of 'duplicate' bifurcates: X and Y are *intrinsic duplicates* iff there is a one-one correspondence between the parts of X and the parts of Y that preserves all *intrinsic* perfectly natural properties and relations; X and Y are *local duplicates* iff there is a one-one correspondence that preserves *all* perfectly natural properties and relations.³⁰ A *local* property is now defined simply as a property that can never differ between local duplicates. That perfectly natural properties are *local* is now built into the definitions.

Now we face a dilemma. I suppose we will want the revised theory to incorporate the Humean denial of necessary connections between distinct existences, in particular, between a point p and its surrounding space: any other point q could have taken p 's place. Of course, at no world is q itself in the place of p . We need to formulate the principle in terms of duplicates: for any points p and q , perhaps from spaces of different worlds, there is a world whose space is a duplicate of the space of p except that it contains a duplicate of q where the duplicate of p would be.³¹ But, on the revised theory, we must decide: do we mean *local* duplicate or *intrinsic* duplicate? Does the principle require that there be a *local* duplicate of q where the duplicate of p would be? It had better not. For example, suppose that p is surrounded by positively curved space, q by negatively curved space. Then, a world whose space is a duplicate of the space of p but with a *local* duplicate of q in p 's place must be both positively curved and negatively curved in the immediate neighborhood of q . No world is like that. So the principle requires only that there be an *intrinsic* duplicate of q where the duplicate of p would be. More generally, the Humean denial of necessary connections is formulated in terms of *intrinsic* duplicates, not *local* duplicates. On the revised theory, however, that will be too weak to capture the spirit of the Humean denial. It rules out necessary connections between the *intrinsic natures* of distinct things. But, on the revised theory, there may be more to a thing than is given by its intrinsic nature. Thus, formulating the Humean denial in terms of intrinsic duplicates fails to rule out necessary connections between the distinct things themselves, in particular, between a point and its surrounding space.³²

I suggest we drop the revised theory and pursue a different tack. Although the local metric, as standardly conceived, is an extrinsic property of points, and therefore not perfectly natural, perhaps the extrinsic local metric is "grounded" on an intrinsic, perfectly natural property of points. To illustrate the sort of grounding I have in mind, consider mass density. If one assumes that each neighborhood of a point has some determinate (finite) mass and volume, then the *mass density* at a point may be characterized as the limit of the ratio of mass to volume, as volume shrinks to zero. So characterized, mass density is an extrinsic property of points. But it is customary in physics, when considering a continuous matter field, to instead take mass density to be a primitive scalar

field: a function that assigns to each point a real number representing (given appropriate units) the *intrinsic mass density* at the point. Given intrinsic mass density, and an assumption about its smooth distribution, mass can be defined by integration. Extrinsic mass density then supervenes upon intrinsic mass density. And, thanks to a fundamental theorem of integral calculus, the values of extrinsic and intrinsic mass density coincide. Note that the smoothness of the intrinsic mass density field is a contingent feature of worlds with continuous matter. Given a principle of recombination for points, there will be worlds whose intrinsic mass densities (perhaps no longer properly so-called) are jumbled up in such a way that no (finite) masses (and no extrinsic mass densities) exist at the world.³³

The suggestion, then, is to say something analogous about the local metric: the extrinsic local metric supervenes on an intrinsic local metric (plus manifold structure). It is the *intrinsic* local metric properties that are perfectly natural. That is on the right track, I think; but there is a problem. Whereas the mass density at a point is a simple *scalar* quantity, the local metric at a point is a *tensor* quantity. How can a tensor be intrinsic to a *point*? Points are spatially simple. Tensors, being operators on vectors/spaces, are spatially complex. It is repugnant to the nature of a point to suppose that a local metric, which is a tensor, could be intrinsic to a point. If we hope to ground the extrinsic local metric on an intrinsic local metric, the latter had better be intrinsic not to a point, but to something spatially complex.³⁴

No sooner said than done. If we are willing to posit perfectly natural properties on theoretical grounds, we should be willing to posit appropriate entities to instantiate those properties: in this case, entities that are spatially complex. I propose that we reify talk of the "infinitesimal neighborhood" of a point. The tangent space at a point is now conceived as the infinitesimal neighborhood of the point "blown large," as viewed through a "microscope" with infinite powers of magnification; it no longer depends for its existence upon the manifold structure. Tensor quantities are intrinsic not to points, but to the infinitesimal neighborhoods of points. At local Gaussian worlds, space (or spacetime) has a "non-standard" structure. There are "standard" points, and there are "non-standard" points that lie an infinitesimal distance from standard points. The points along a path in space are ordered like the non-standard continuum of Abraham Robinson's non-standard analysis.³⁵

Let us take stock. The local Gaussian conception of distance, if founded upon standard differential geometry, is committed to local metric properties that are both extrinsic and perfectly natural. I propose founding the local Gaussian conception instead upon non-standard differential geometry. That allows the perfectly natural local metric properties to be intrinsic, though not to points, but to their infinitesimal neighborhoods. The intrinsic local metric at a point now comprises a family of infinitesimal distance relations. So there turn out to be perfectly natural distance relations after all; but they are local, not global, because they hold only among points within an infinitesimal neighborhood of a standard point. ('Local' is defined with respect to the topology of the

standard points.) The local Gaussian is no longer committed to perfectly natural, extrinsic properties. Metaphysical worries about necessary connections have been resolved.

When non-standard analysis gained mathematical and logical respectability some thirty-odd years ago, the question naturally arose whether the non-standard continuum is instantiated at any possible world, or even at the actual world. Perhaps the mere consistency of non-standard analysis already gives reason to suppose that the non-standard continuum is possibly instantiated. The role that non-standard differential geometry can play in firming up the metaphysical foundations of physical theory gives reason all the more—including reason to suppose the non-standard continuum is actual.

NOTES

1. What if quantum mysteries are taken into account? Should I endorse some conception of space as foamy, or spongy, or stringy, or loopy? I haven't a clue.
2. For the conventionalist line, see Hans Reichenbach, *The Philosophy of Space and Time* (New York, 1957), or Adolf Grünbaum, *Philosophical Problems of Space and Time*, 2nd ed. (New York, 1963), chapter 16. For a realist response, see Graham Nerlich, "Is Curvature Intrinsic to Physical Space?" *Philosophy of Science* 46 (1979): 439–58.
3. For a discussion of which spatial structures are possible, that is, instantiated at some possible world, see my "Plentitude of Possible Structures," *Journal of Philosophy* 88 (Nov. 1991): 607–19.
4. For convenience, I include ∞ among the non-negative reals, and read ' $D_\infty(\mathbf{p}, \mathbf{q})$ ' as \mathbf{p} stands in no distance relation to \mathbf{q} . As usual, ∞ is greater than any other non-negative real, and ∞ added to any non-negative real is ∞ . This artifice allows the distance axioms to hold at worlds with "island universes," spatially disconnected parts.
5. Those who prefer genuine transworld identity may simply suppose that one of the appropriate counterpart relations is the relation of identity. No trouble arises because, for the counterpart relations introduced below, counterparts are always intrinsic duplicates.
6. Following David Lewis, for discussion, see *On the Plurality of Worlds* (Oxford, 1986), 59–61.
7. Again following David Lewis, *On the Plurality of Worlds*, 61–62. The definitions below of 'intrinsic', 'internal', and 'external' are also adapted from Lewis, although he does not apply the word 'intrinsic' to relations. Note that quantifiers here and below range over all *possibilia*.
8. As Lewis notes (*On the Plurality of Worlds*, 68), on a theory of universals according to which universals are parts of the particulars that instantiate them, we must everywhere replace the fusion $X + Y$ with an *augmented fusion*, which includes among its parts not only X and Y , but the dyadic universals that hold between X and Y (or between parts of X and parts of Y), and the monadic universals that hold of the fusion $X + Y$ (or of its parts). *Mutatis mutandis* for a theory of tropes.
9. The intrinsic relations can be further divided into *internal* and *external*. A (dyadic) relation is internal just in case whether it holds of a pair $\langle X_1, X_2 \rangle$ depends only upon the intrinsic nature of X_1 and of X_2 , not upon the intrinsic nature of $X_1 + X_2$. In terms of duplicates: A (dyadic) relation R is *internal* iff, for all X_1, X_2, Y_1, Y_2 , if X_1 and Y_1 are duplicates and X_2 and Y_2 are duplicates, then R holds $\langle X_1, X_2 \rangle$ iff R holds of $\langle Y_1, Y_2 \rangle$. R is *external* iff R is intrinsic but not internal. This distinction, however, will not be needed below. It is agreed on all sides that distance relations, if intrinsic, are external.
10. On the interpretation of *de re* counterfactuals, see David Lewis, *Counterfactuals* (Cambridge, Mass., 1973), 36–43.

11. That is, the points of A form an open set in the usual topology. A part of space is open iff it excludes all of its boundary points.
12. The intrinsic conception of distance is endorsed by David Lewis in *On the Plurality of Worlds*, 62. Lewis does not consider alternative conceptions.
13. What about a right- and a left-handed glove that are mirror images of one another, and thus congruent? That is no counterexample. The gloves differ in *orientation*, and orientation is not intrinsic, as consideration of a Möbius strip (or its three-dimensional analog) should make clear.
14. Of course, a further generalization is needed to account for the interval relations of Minkowski spacetime, since the interval squared may be positive, negative, or zero.
15. In what follows, unless otherwise noted, I assume that paths are continuous and "smooth"—that is, without corners or cusps. (Technically, I assume that paths can be given a parameterization that is differentiable with non-zero derivative at all points along the path.) In order that the notions of continuity and "smoothness" be applicable to parts of space, the Gaussian must assume that space is a *manifold*, that space has both topological and differential structure.
16. I assume island universes are possible. An upholder of the Gaussian conception who denied this should refuse to answer question (1) on the grounds that there is no world whose entire space is a duplicate of X_1 .
17. This is because the surface of a sphere (unlike a flat plane) is *rigid* in E : any part of E that is isometric to the surface of a sphere in E is also congruent to it, and so itself the surface of a sphere of the same size. Intuitively, a surface in E is rigid if it cannot be deformed without stretching or tearing.
18. Note that on the intrinsic conception, \mathbf{p}' and \mathbf{q}' are separated by more than a quarter great circle of Y_6 , since \mathbf{p}' and \mathbf{q}' are twenty feet apart in Y_6 , and Y_6 is the surface of a sphere of circumference eighty feet.
19. For an excellent introduction to differential geometry, including the standard motivational asides, see Barrett O'Neill, *Elementary Differential Geometry* (New York, 1966).
20. Unfortunately for our purposes, distance-within-a-surface is often called "intrinsic distance," since it is part of the surface's intrinsic geometry. I will avoid that usage.
21. This notion of a boundary point of space requires differential (but not metrical) structure: a boundary point is one such that some path to the point cannot be "smoothly" extended.
22. *Globally* supervene, that is, since the assumption only applies to entire worlds. On the distinction between various notions of supervenience, see Paul Teller, "A Poor Man's Guide to Supervenience and Determination," *Southern Journal of Philosophy* 22 (1984): 137–62.
23. More exactly, the metric tensor at a point is an inner product on the tangent space of the point; and the metric tensor field is differentiable, it varies "smoothly" from point to point.
24. The information carried by ds^2 is coordinate-independent, though of course calculations of length of path will be done by representing ds^2 and the path in question relative to some chosen coordinates. (Coordinate-free geometric objects are represented in boldface.) In the case of a three-dimensional Euclidean space, there will be x, y, z -coordinates under which $ds^2 = dx^2 + dy^2 + dz^2$. In general, however, ds^2 will be a more complicated quadratic function of dx, dy , and dz , for any x, y, z -coordinates.
25. For an elaboration on this theme, see the introductory chapter of Charles Misner, Kip Thorne, and John Wheeler, *Gravitation* (San Francisco, 1973), 4.
26. David Lewis, Introduction to *Philosophical Papers*, Volume II (Oxford, 1986), 3–4.
27. David Lewis, in the passage just-quoted, requires that local properties be intrinsic properties of points (or their point-sized occupants); but I do not think that is how 'local' is standardly used in mathematics or physics.
28. Neighborhood-dependent properties may be *exclusive* or *inclusive*: those are exclusive

that exclude information about the intrinsic nature of points that instantiate them. Thus, a property P of points is *exclusively neighborhood-dependent* iff, for any points p and q , for any neighborhood N of p and any neighborhood M of q , if $N - p$ is a duplicate of $M - q$, then P holds of p iff P holds of q .

29. This is a bit fast and loose. Unless the paths through p are embedded in some higher-dimensional Euclidean space, the derivatives in question are not defined, and tangent vectors are instead identified with (directional) derivative operators. (See O'Neill, *Elementary Differential Geometry*, 182–84.) But the argument is essentially unchanged, since derivative operators, which require manifold structure, are no less neighborhood-dependent than derivatives.

30. One might ask, independently of the question whether perfectly natural properties can be extrinsic, whether 'duplicate' in ordinary usage means 'local duplicate' or 'intrinsic duplicate', or is indeterminate. *Test case*. Consider a cube with sides of two feet and a sphere with a diameter of one foot, each composed of (the same kind of) homogeneous continuous matter. The sphere has continuum-many intrinsic duplicates among the parts of the cube; but the sphere has no local duplicates, since no interior point of the cube is a local duplicate of any boundary point of the sphere. Using our ordinary notion of duplicate, how many duplicates of the sphere are there in the cube? It seems to me one can answer either way.

31. This is an instance of the principle of recombination put forth by David Lewis. See *On the Plurality of Worlds*, 86–92. (Of course, the points in question must be of the same kind, be it Newtonian or spatiotemporal.)

32. The argument is especially compelling if one holds that perfectly natural properties correspond to immanent universals or classes of tropes. For immanent universals or tropes are present in their instances. Now consider a neighborhood-dependent, perfectly natural property of p . The corresponding universal, or a corresponding trope, is present at p . And, unlike a dyadic universal or trope, it is wholly present at p . (Remember: its holding at p tells one nothing about any point other than p , not even something relational.) Intrinsic or not, how can one deny that it is part of the nature of p , and so must be "recombined" along with p ?

33. Michael Tooley has argued that extrinsic (or "Russellian") velocity should be grounded in this way on primitive velocities that are intrinsic to points. But the theoretical reasons for positing primitive velocities, at least at worlds approximating ours, seem to me much weaker than the theoretical reasons for positing primitive local metrics. See Michael Tooley, "In Defense of the Existence of States of Motion," *Philosophical Topics* 16 (Spring 1988): 225–54.

34. Denis Robinson asks whether vectors could be intrinsic to points, and answers "no," in "Matter, Motion, and Humean Supervenience," *Australasian Journal of Philosophy* 67 (December 1989): 394–409. I concur. Although vectors are spatially less complex than tensors, they have a "tail" and a "tip": too much to fit within a single point.

35. Non-standard analysis is applied to differential geometry, for example, in Abraham Robinson, *Non-Standard Analysis*, rev. ed. (Amsterdam, 1974).