

Relative Locations

Andrew Bacon*

October 13, 2016

Abstract

The fact that physical laws often admit certain kinds of space-time symmetries is often thought to be problematic for substantivalism — the view that space-time is as real as the objects it contains. The most prominent alternative, relationism, avoids these problems but at the cost of giving abstract objects (rather than space-time points) a pivotal role in the fundamental metaphysics. This incurs related problems concerning the relation of the physical to the mathematical. In this paper I will present a version of substantivalism that respects Leibnizian theses about space-time symmetries, and argue that it is superior to both relationism and the more orthodox form of substantivalism.

Substantivalism is the view that locations exist independently of the objects that they are locations of. Thus, for example, the moon’s present location would have existed even if it had not been located there and, reciprocally, there are regions that could have been the location of some material object, but aren’t (unoccupied regions).

Substantivalism can have a number of different motivations. In ordinary English we frequently talk about, and quantify over, locations. For example, I can talk about where Emily went on vacation, or where I left my keys; I can in some cases talk about unoccupied locations, such as the place I would have been had I traveled halfway to the moon. The substantivalist has a ready answer for what we are doing when we quantify in this way: we are quantifying over regions of space-time. If this were the central motivation, however, it would suggest that the dispute is primarily about the existence of locations. This is not quite right: even those who oppose substantivalism — the relationists — are often happy to accept ordinary talk about locations provided that that talk can be recovered from a more acceptable fundamental ontology. (For these theorists, locations are a bit like shadows and holes: they are, metaphysically speaking, ‘second-rate’ entities in some sense. We can readily quantify over them, but their existence depends on other kinds of things: a location on the thing that occupies it¹, much like a shadow on the thing that casts it, and a hole on the thing it penetrates.)

Substantivalists arguably should not accept this reduction of our ordinary talk of locations to quantification over regions of space-time either. My pocket is not a region of space-time, and nor is Paris — for example, my pocket changes its shape as it goes through the washing machine, and a region of space-time cannot change its shape. But my pocket

*Many thanks to Jim van Cleve, John Hawthorne and Jeff Russell for invaluable discussion and to an audience at Toronto who saw an earlier version of this paper. Thanks also to the judges Brad Skow, Jessica Wilson and Cian Dorr and the editor Karen Bennett for many helpful suggestions.

¹Or perhaps, if the region is empty, the possible objects that could have occupied it. See Forbes [14].

and Paris are the kinds of things I might be asking about when I ask where my keys are or when I ask where Emily when on vacation.²

Since I find the above sorts of objections compelling, I shall place little weight in what follows on the arguments from ordinary language. In my view, the strongest motivations for substantivalism — the motivations that play a central role in this paper — are slightly more theoretical. Space-time points tend to appear in the formulation of many of our simplest physical theories and this gives us good reason to take their existence seriously.

Unfortunately the most straightforward versions of substantivalism and relationism suffer from having a pair of undesirable, and arguably related, consequences. Substantivalism predicts the existence of physically indistinguishable worlds that differ solely concerning where things are located. For according to substantivalism there could be two worlds that differ from each other only in that each object's location in one world has been displaced relative to the other by some fixed distance in some particular direction (see Leibniz [18]).

By refusing to take talk of locations and regions of space-time at face value the relationist does not face this problem in its most acute form. However relationists are subject to a somewhat similar concern. The most straightforward versions of relationism take geometrical properties and relations to be a matter of fundamental relations holding between material objects and abstract objects, such as numbers to represent distances between objects, vectors to represent forces, and so on. However much like the substantivalist's location relation, I shall argue, one can have physically indistinguishable worlds in which the relations between the material objects and the appropriate abstracta have undergone a similar kind of displacement.

In this paper I shall in response to these worries be advocating for a novel version of substantivalism underwritten by a non-standard account of the relation between material objects and the regions they occupy. As I cautioned against earlier, the theory is not intended to model our ordinary way of speaking about locations. However the theory does allow one to give a story about how the shapes of objects, the distances between them, and other such facts can emerge from the relations they stand in to a sufficiently structured space-time manifold.

Here is an outline of the paper. In section 1 I review Leibniz's shift argument against substantivalism and argue that it provides us with reasons to look for alternatives to the orthodox version of substantivalism. In section 2.1 I consider the analogous problem involving abstracta for (versions of) relationism, and consider the options for a relationist who wishes to give an account of geometrical properties without reference to abstract objects. To this effect I outline in section 2.2 a result that says that any first or second-order language that can express, relative to some class of models, all possible configurations of a system of n particles either must have an infinite set of fundamental primitives or must admit a model with an infinite ontology (whether that be an ontology of space-time points, abstract objects, or something else). In section 3 I develop a substantivalist theory according to which an object is located at a region of space-time if and only if it is located at every transformation of that region. Thus unlike orthodox substantivalism, objects are multiply located in a fairly far-reaching way. Finally, I consider two further technical issues in the appendices: the treatment of space-time fields and the first-order theory governing the interaction of parthood and location.

²Of course some people — the *supersubstantivalists* — do identify my pocket with a region; however some fairly elaborate maneuvers are needed to make sense of the idea that ordinary things could have had different shapes.

1 Leibniz’s shift argument

Substantivalism is subject to an old objection that traces all the way back to Leibniz [18]. Pick some direction and distance, and imagine a world at which the location of every object has been translated uniformly in that direction by that distance, but in which every other property of the object has been kept the same. According to the substantivalist, worlds related by such operations are genuinely different because they disagree about facts concerning which particular regions of space-time each object is located at. Substantivalism thus falls afoul of the following seemingly attractive thesis:

NO SHIFTS: There are no differences between shifted worlds.

A pair of worlds are ‘shifts’ of one another if they are related by a translation of every object’s location of the sort described above. The crucial feature of shifts to highlight here is that they preserve all the laws of physics, and arguably all observable properties and relations between objects.³ Very similar arguments could be made by appealing to other spatial transformations operations: rotating about some axis or reflection about some plane. And just as we can consider spatial transformations, we may also consider operations involving time: time reversal (a kind of temporal reflection), temporal shifts (moving every event forwards or backwards a fixed amount of time) and boosts (uniform shifts of velocity), and of course, arbitrary combinations of any of these operations. We shall call a combination of these operations a Galilean transformation.⁴

The principle NO SHIFTS has been given a number of distinct motivations over the years, some less convincing than others. While some have overtly theistic or verificationist premises, NO SHIFTS has remained central to the debate about the existence of space-time, even though attempts to motivate it from more general principles have changed. Most contemporary philosophers take the best justification for NO SHIFTS to be a defeasible one: theories that postulate undetectable structure that play no role in explaining the observable world are not to be preferred over theories that do not postulate this structure but are otherwise just as simple and explanatory (see, for example, Russell [25], Pooley [24]) As a silly example, one could imagine a theory which postulates the existence of an *absolute origin* from which every point can be assigned an absolute distance. Although a particular origin could well make a nominal appearance when the dynamical laws are formulated in terms of a particular coordinate system, a world in which an alternative point were the absolute origin would be identical to our own in all physical respects. This kind of additional structure is completely undetectable and is also unnecessary to the formulation the dynamical laws. A less silly example is Newton’s original formulation of classical mechanics, in which there was a distinguished velocity called ‘absolute rest’: objects traveling at that velocity counted as being at absolute rest, and objects moving relative to it had absolute motion. The move

³One might object here on the grounds that one can have singular thoughts about particular regions of space-time, and thus knowledge, that isn’t preserved under transformations. (I can know that I’m located right here, at this particular region of space-time for example — I clearly wouldn’t know this in a world where I wasn’t located right here). I’m going to set this issue to one side for the time being, but see Maudlin [19] for more discussion.

⁴Note on terminology: sometimes reflections of various sorts are excluded from the definition of a Galilean transformations — I always intend them to be included in what follows. Under the operation of composition these transformations form a group called the Galilean group. Note, however, that I do not include enlargements in the space of transformations: these are not in fact symmetries of the laws of Newtonian physics, since the rate at which two particles of equal mass will accelerate towards one another depends on the distance they are separated.

from Newtonian space-time to *Galilean space-time*, in which this redundant structure is eliminated, is now universally accepted.

The main problem for this line of reasoning is that unlike the case of the absolute origin and absolute rest, we don't have a clear alternative theory in which the redundant structure is eliminated. Thus, while many philosophers take this version of Leibniz's challenge to be serious one, ultimately reject it on the grounds that the most prominent alternative to substantivalism — relationism — doesn't share the theoretical virtues of the substantivalist theory (see for example Pooley [24] and Maudlin [21] p65-66).⁵ This strikes me as sufficient reason to explore alternative theories which don't posit this sort of redundant invisible structure.⁶

2 Relationism

The best argument for NO SHIFTS rests on general considerations of theory choice, and is thus successful only if there is a simple alternative to the standard version of substantivalism that respects it. It is often assumed that relationism is this alternative. Unfortunately relationism faces a somewhat similar set of issues.

2.1 Mathematical Relationism

According to relationism, the motions and geometrical properties of material objects are not grounded by the relations they stand in to a background space-time manifold. Thus the relationist is in need of some other kind of theory of these sorts of properties and relations. To keep things simple we shall restrict our attention to a relatively sparse world consisting of three point particles, arranged in a certain configuration, traveling with certain momenta as governed by the laws of Newtonian gravitation. Since it is clear that there is a difference between worlds in which the particles are arranged in, say, a regular triangular shape, and those in which they are colinear, or worlds in which the masses of the particles, and thus the forces between them, are greater or smaller, the relationist owes us an account of what these differences consist in.

According to the most straightforward versions of relationism, geometrical properties like these consist in fundamental relations holding, in certain configurations, between the material objects, numbers and other kinds of abstract objects. Thus, for example, distances can be represented by a fundamental three place relation, $Dxyz$, taking two particles and a real number, satisfying the axioms of a metric space. (In mathematical contexts, it is standard to equivalently describe such a relation as a function $d(x, y)$ mapping x and y to the distance between them.) Thus three particles arranged in a regular triangle are such that each pair of them is related to the same number. The gravitational force exerted by one particle on another can similarly be represented by a relation, $Fxyz$, taking two particles and a vector from some suitable normed three-dimensional vector space, and mass can be


⁵Although there have been several attempts to develop a 'Machian' theory of this sort, it's unclear how successful they are. See Barbour and Bertotti [4] and Barbour and Bertotti [5]. See Pooley [24] for an overview.

⁶This general sort of project has been attempted by Dasgupta ([8], [REF]) and Russell [25], but both sorts of theories involve contentious metaphysical ideology such as a many-many grounding relation or a primitive factuality operator. The approach I will be recommending here is straightforwardly intelligible to anyone who has the concepts presupposed in the original relationism/substantivalism debate — specifically, anyone willing to theorize in terms of the *location relation*.

represented by a binary relation between particles and real numbers Mxy (equivalently a function $f(x, y)$ mapping pairs of particles to vectors and a function $m(x)$ mapping particles to real numbers).⁷ Given this kind of background theory, one can begin to attempt to reconstruct a relationistically acceptable version of Newtonian physics: for example, we might subject F the restriction that if F holds between x, y and v it holds between y, x and $-v$ according to Newton's third law.⁸ Although there is room for much variation in the details let us call this general kind of approach, in which geometric and physical properties are grounded in relations between the physical and the platonic realms, *mathematical relationism*. (For a prominent theory of this sort, which freely employs relations to mathematical objects, see Barbour and Bertotti [5]).

The framework just outlined seems to be susceptible to an objection similar to the one afflicting substantivalism: it seems that one can describe two different configurations of these relations between objects and numbers that correspond to the same physical scenario. For consider any legal arrangement of our three particle system: this will consist in the relations D, F and M holding between our particles, numbers, and vectors in some configuration. In particular, each pair of particles, x and y , will be related by F to some vector v in an abstract vector space V , representing the force that x exerts on y . Now consider another configuration which agrees with the original regarding which things are related via M and D , but in which the vector related to each pair of particles by F has been uniformly switched, at every time, for the result of rotating that vector by 90 degrees about some fixed axis in the abstract vector space it belongs to. Notice that this operation will preserve everything of importance for the modeling of forces — crucially, it will preserve the inner product which represents the lengths of the vectors and the angles between them.

This operation bears many of the marks of Leibniz's shift argument: the operation of rotating the vector that represents the force between any pair of particles appears to be a symmetry of our relationist theory. Although this is not on its own a sufficient reason to reject this theory, it does mean it lacks the advertised advantages over substantivalism. We must, however, be careful to distinguish the symmetry just described from a different operation. For example, one can imagine a possible world that is not legal by the standards of the Newtonian theory of gravitation, in which the force that the particle y exerts on x is not parallel, but orthogonal to the line passing between them — see the 'Illegal world' diagram below. (Note: each diagram is supposed to represent three different states all at the same time t).



This is not the operation I am describing: if the force y exerts on x really is acting perpendicular to the line between them then at later times x would accelerate in a direction

⁷I have set aside the issue of time here. A proper treatment may involve introducing another argument place to these predicates for a time, or giving them a tense-logical treatment (the latter seems natural for those relationists, such as Arthur Prior, who reject time as well as space).

⁸One might wonder in one could eliminate forces just by appealing to distances and masses (or perhaps masses in favour of distances and forces). This is far from clear: for example a pair of equally massive particles orbiting one another in a circular motion will have constant masses, and be at a constant distance from one another, although the force acting between them is constantly changing, so that forces can't be recovered from mass and distance alone.

perpendicular to that line, so that the distance between x and y ought to increase for a period of time (absent other forces).⁹ In the original example, however, I stipulated that the distance facts are the same at both possibilities at all times, so if the original case is legal, the distance between x and y will decrease for a period in the variant world. The ‘shifted’ world is therefore a world in which the physical force in fact points from x to y (since that is the direction that x will move in), but in which its so pointing consists in x and y standing, via F , to a different abstract vector (u instead of v , as in ‘Legal world 1’ and ‘Legal world 2’).

Of course, one could insist that one of these configurations is not really a metaphysical possibility. Perhaps, in that configuration, only v but not u could possibly represent the force exerted by y on x . Similar moves can also be made in the substantival case: perhaps no two worlds that agree about the non-locational profile of the material objects can disagree about their locations in space-time. But such responses do little to assuage the feeling of arbitrariness — why, for example, is it necessary that when the particles are arranged this particular way, they are always located at *this* region of space-time.

Note also that the vector representing the force y exerts on x can differ between *times*, even when those times agree about the mass and distance facts. For consider a system of two equally massive point particles, x and y , orbiting one another in a circular motion. The distance between the particles, and their masses remain fixed over time. However the vector representing the force between x and y is rotating at a constant rate.¹⁰ Thus, in fact, the states depicted in ‘Legal world 1’ and ‘Legal world 2’ could represent the state of the same world at two different times. So the thought that it’s impossible for v or impossible for u to represent the force acting between x and y when the distance and mass facts are in such a state seems to rule out certain kinds of physically possible scenarios.¹¹

It may strike the reader that the above problem is specific to the particular form of mathematical relationism I have outlined above, which relied on the use of vectors. However the problem outlined above is but an instance of a more general problem for relationism outlined by Hartry Field in [11]. For example, Field notes that we can describe the world equally well using different units (see Field [11]). Thus a configuration of M , D and F in which M relates each particle to its mass in kilograms, and the configuration in which M relates each particle to its mass in pounds but is otherwise the same, correspond to the same physical scenario. The ‘shift’ that we are performing in this case is that of multiplication by a factor of a scalar quantity, rather than a rotation of a vector quantity.¹² There is something

⁹We are assuming here that the world is legal in at least the sense that particles accelerate in the direction of the forces acting on them, even though it is illegal in the sense that the force exerted by a particle does not point towards the particle.

¹⁰This is part of the reason why this version of relationism can avoid Newton’s bucket.

¹¹One could try to make this argument more rigorous by appeal to the principle that whatever sometimes happens could happen (see Dorr and Goodman [9]), for in the two particle world described, x and y sometimes bear F to v and sometimes to u , but always bear M and D to fixed numbers. (Although I do not myself subscribe to this principle, there is something very compelling about the intuition in this case.)

¹²In the case of scale dependence there’s a fairly straightforward fix: instead of having a primitive binary relation Mxy relating each object to a number, have a ternary relation $Mxyz$ relating two physical objects to a number that simply tells us what the *ratio* of the mass of x and y is. Whatever units we choose, the ratio of two masses will remain the same. However this doesn’t really speak to the underlying problem that when abstract numbers appear in physics there are usually other abstract objects that would do the job just as well; usually the choice of mathematical objects to use is made based on convenience. A slightly contrived version of this problem applies even to ratios of physical quantities — there are many abstract objects that could do the job of ratios equally well. Ratios are real numbers, and set theoretically we can construct these entities in several equally natural ways: we can represent them by certain sets of rational numbers, called Dedekind cuts, or by certain kinds of converging sequences of rational numbers called Cauchy sequences.

very suspect about a fundamental metaphysics in which one has to choose between kilograms and pounds, quite aside from its the fact that it generates invisible distinctions.

The issue to do with units above appears to be part of a broader set of problems afflicting theories that postulate relations to mathematical objects in order to account for the fundamental properties of physical objects. In general there is a lot of arbitrariness in our choices when we represent the world using abstracta: not just in the choice of units, but sometimes in more far reaching ways. For example, Newtonian mechanics is simple enough to be formulated in terms of affine spaces — in which pairs of space-time points are related to objects in a mathematical vector space — but it could also be formulated in the language of differential geometry in which case each space-time point is instead related to an equivalence class of real-valued smooth functions on a neighborhood of that point. Similar points extend to most uses of abstracta by physicists, whether it concerns the choice of origin and orientation of a coordinate system, the choice of units, or sometimes the very choice of mathematical formalism itself.¹³

Thus some philosophers, such as Field, have maintained on this basis that ‘relations between physical things and numbers are conventional relations that are derivative from more basic relations that hold among physical things alone’. Field has in mind relations of magnitude in particular and, although he is presumably at least partly being motivated by his nominalism, also mentions considerations such those about units mentioned above. I am not a nominalist, but I think the idea is compelling nonetheless. Even if one rejects nominalism, it would be puzzling if the fundamental laws of motion, for example, depended for their truth on the existence of abstract objects. One could dramatize this intuition by imagining that all numbers were to disappear tomorrow: insofar as we can entertain such a hypothesis, it doesn’t seem likely that there would be any serious consequences for the non-mathematical universe — it’s not like the earth would stop orbiting the sun, or that planes would start falling out of the sky. Gravity, for example, seems to be a physical force whose existence does not depend on the numbers we use to represent it: surely concrete things could have moved about in the way that physics demands even if there hadn’t been any numbers.

2.2 The prospects for alternative versions of relationism

It is best, I think, for a relationist to maintain that material objects have their topological, metrical and geometrical structure *intrinsically*: structure that is internal to the domain of material objects and is not merely inherited by their relation to a platonic realm of numbers. There are a couple of ways we could go about this project, depending on whether we take there to be an infinite number of metaphysically primitive geometrical properties or a finite number:

1. GEOMETRIC PRIMITIVISM: Take each geometrical property and relation to be metaphysically primitive. On this picture all geometrical properties are equally fundamental.

This theory will have an infinite number of primitives: for each possible value of α there is a primitive relation between particles of being α meters apart, for each possible

Even if real numbers are *sui generis* entities, and are not identical to either of these constructions out of rational numbers, that just means that we have one more isomorphic mathematical structure to choose between.

¹³For example, some physicists prefer to theorize in the language of category theory, which often leads to certain kinds of more familiar abstracta being replaced with categories.

shape, there is a primitive fundamental property of having that shape, and so on.

2. GEOMETRIC REDUCTIONISM: Attempt to fix the geometrical structure by a smaller finite set of primitives.

Perhaps to the notion of an object being an *open sphere* (Tarski [30]), or to the notion of two pairs of particles being *congruent* to another, and one being *between* the other (Hilbert [15], Tarski [29]).¹⁴

If one thought that our fundamental properties and relations are governed by simple laws, the first option should strike us as deeply unsatisfactory. There are, for example, general geometrical laws relating collections of particles with certain shapes, and moreover physical laws governing how those shapes should evolve as the particles are attracted to and repel one another. In the first kind of theory we have no hope of writing these sorts of laws down: to achieve quantification over distances (which we need to do in order to talk about rates of change, for example) would require one to employ large infinitary conjunctions and disjunctions.¹⁵

The latter sort of approach to geometry has been developed extensively by Hilbert and Tarski, and has most prominently been championed in the philosophical literature by Field [11] (see also Arntzenius and Dorr [3], Casati and Varzi [7], Maudlin[20], for some similar approaches to geometrical and related structure). To illustrate let us focus on the work of Hilbert [15] and Tarski [29]. The usual way to represent distances between points would be to introduce a metric — a function from pairs of space-time points to real numbers — telling us how far apart the points are. The mark of the Hilbert-Tarski program, by contrast, is that it employs a small number of geometrical properties and relations whose relata are concreta. These geometrical properties can be understood in both a substantialist or a relationist setting by interpreting the primitives as applying to either regions of space-time or material objects respectively. In Tarski’s geometry of solids, for example, we have the following primitives:

1. A binary relation $x \leq y$ whose intended interpretation says that x is a mereological part of y .
2. A unary predicate, Sx , stating that x is an open sphere.¹⁶

Hilbert’s axiomatisation, also later refined by Tarski, instead only invokes a pair of relations whose arguments are point-like objects:

1. A three place betweenness relation, $Bxyz$, whose intended interpretation states that x lies on the straight line segment between y and z .
2. A four place spatial congruence relation, $Cxyzw$, whose intended interpretation states that the distance between x and y is the same as that between z and w . (For short: xy is congruent to zw .)

¹⁴If we are interested in properties relating to the smoothness of the manifold the notion of a converging sequence of points (Dorr and Arntzenius [3]), or topological properties the notion of two things touching (Casati and Varzi [7]) or a closed line (Maudlin [20]).

¹⁵For a comprehensive discussion of related issues see Skow [28]; see also the further discussion in Kleinschmidt [17].

¹⁶This is what modern mathematicians would call an *open ball*; the word ‘sphere’ is now typically reserved for the two dimensional surface of a ball, however this is not what Tarski meant by a sphere.

It should be noted that in this setting the quantifiers are understood as ranging over point-like objects (space-time points or point-like material objects); if one wanted to make that restriction explicit one could introduce another primitive applying to mereological atoms, or one could simply stipulatively understand congruence so that it applies to no complex objects and define atomicity as standing in congruence relations to some things.

The general approach is not limited to quantities representing distances either. Field has shown that the latter sort of theory can be extended to other quantities. If I want to talk about the numerical value of a field at a given point — for example, the mass density field or the gravitational potential — I can employ a similar trick. One can introduce a congruence relation stating that the difference in gravitational potential, for example, at x and y is the same as the difference of the gravitational potential at z and w , and a betweenness relation saying the potential at x is between the potential at y and z respectively (see Field [11]).¹⁷

One might hope that a relationist could employ small set of geometrical primitives, such as those above, and attempt to recover geometric structure that way.

Unfortunately, as Hartry Field has noted in [12], some of the above theories are simply not available to the relationist. In particular the Hilbert-style theory, employing congruence and betweenness, fix the relevant geometric structure only if there is a sufficiently large number of geometric objects hanging around. This can be illustrated with a simple example: consider a situation in which the distance between x and y is two times the distance between z and w . While this may on the surface look as though we have a relation between four points and the number two, in the Hilbertian setting reference to the number two can be eliminated. This relation can be stated instead as follows:

There is some point u between x and y such that both xu and uy are congruent to zw .

If we wanted to say the distance between x and y was three times the distance between zw we'd say there was a u_1 and u_2 such that xu_1 , u_1u_2 and u_2y were each congruent to zw . One can see, without much trouble, how to paraphrase away talk of arbitrary rational ratios of distances in this fashion. Note that once you have pinned down the rational distances between points all remaining distances between points are fixed. So, given the existence of enough point-like objects, betweenness and congruence facts are enough to pin down all metric structure: no two worlds can agree about the betweenness and congruence facts and disagree about the geometrical facts.

The above statement captures the notion of a pair of particles being twice as far apart as another pair perfectly well in a substantialist setting, since according to that theory, for any two space-time points there is another space-time point between them. However Field notes that this is not so for point-sized material objects: if I take the closest point to the earth on the edge of the moon, p , and take the closest point to the moon on the edge of the earth, q , (suppose for a moment that both have definite boundaries) then, at least by the relationist's lights, there are no entities between p and q , not even a space-time point. Thus, even if I were twice as far away from the earth as the moon is, I wouldn't count as such by the lights of Hilbert's analysis given a relationist ontology. Indeed, it is easy to see that there are distinct arrangements of me the moon and the earth that agree about all betweenness and congruence facts but are nonetheless very different geometrically.

¹⁷In order for this to work there must be a rich enough variety of field comparisons for us to be able reconstruct the field values (up to a scale) from the comparisons. This is guaranteed if we make the assumption, typical in physics, that physical fields are always continuous, so that a field is either constant everywhere or inhabits an open interval of the space of possible field values.

The problem Field has identified here seems to be much more general. For example, Tarski defines the topological notion of two spheres touching (i.e. touching but not overlapping) in terms of the notion of a sphere as follows:

x touches y if and only if, x is disjoint from y , and any two spheres containing x and disjoint from y are such that one is part of the other.

Neither this definition nor any other will work in a relationist setting. Imagine two worlds containing two perfectly spherical balls, x and y , in an otherwise empty space. Suppose also that in the first they are touching, and in the second they are not. Both worlds agree with one another concerning which things are spheres and which things are parts of what, but they disagree about which things are touching. Thus, given a relationist ontology, the touching facts are not fixed by the sphere and parthood facts (in particular, because the only sphere that contains x is x itself, according to the relationist ontology, x and y count as ‘touching’ in both worlds by Tarski’s definition).

It follows that in order to pin down the geometrical structure the relationist needs to introduce *twice the distance* and *touching* as new primitives. It is natural to wonder whether this can be done by adding only finitely many primitives; thus evading the undesirable aspects of PRIMITIVISM. To address this question let us focus on a simple example world consisting of only three point particles, arranged in some shape. There are an uncountable infinity of arrangements like this that differ regarding the ratios of the distances between the three particles.

In this setting the primitivist strategy involves introducing, by brute force, an uncountable infinity of primitive ternary relations, so that for each possible arrangement a of the particles x, y, z there is a fundamental relation R_a that holds between x, y, z (in any order) iff those particles are in that arrangement.

It is natural to wonder whether we can do better. Can we specify a theory of three point particles with finitely many primitives without expanding our ontology to include space-time points or numbers? In order to answer this question we need to state it a bit more precisely. An *arrangement* of three particles can be represented by the three particles x, y, z and the ratios of the distances between any pair of them. Formally the set of possible arrangements of three particles, \mathcal{A} , is the set of metric spaces (M, d) such that $M = \{x, y, z\}$ quotiented out by scale.¹⁸ Someone hoping to write down a theory capturing the geometry of three point particles must choose a language \mathcal{L} — given by specifying some set of non-logical primitives — and present a theory which can be either specified axiomatically, or by a class of intended models, \mathcal{C} . For example, in the Hilbert-style theory the primitives of our language are the congruence and betweenness relations, three constants denoting the three particles, and a location relation for stating the location of each particle. The class of intended models consists of three-dimensional Euclidean spaces with the three particles x, y and z located at three points of that space. Each model of the theory ought to correspond to some arrangement of the particles: there ought to be a surjective function $Arr : \mathcal{C} \rightarrow \mathcal{A}$. In the Hilbert-style theory Arr is easy to specify — each model easily determines an arrangement because the underlying metric of the Euclidean space tells us what the distances between the three distinguished points are.

The Hilbert-style theory has a nice feature. For any two models M and M' in \mathcal{C} of the theory in which the arrangement of x, y and z differ (i.e. $Arr(M) \neq Arr(M')$), one can find a sentence in the language of congruence and betweenness, ϕ , such that $M \models \phi$ and

¹⁸That is to say, an arrangement is an equivalence class of such metrics where (M, d) and (M', d') are equivalent iff $d = \alpha \cdot d'$ for some $\alpha \in \mathbb{R}$.

$M' \not\models \phi$.¹⁹ In this way any two arrangements can be distinguished by some sentence of the language.

The primitivist relationist theory also satisfies this constraint. The language of this theory has a primitive relation R_axyz for each possible arrangement $a \in \mathcal{A}$. The models of this theory have a minimal relationist ontology: the domain of each model contains only the three particles x, y, z . Moreover, for any arrangement $a \in \mathcal{A}$ there's some model in which the objects are arranged that way: a model in which R_a applies to x, y and z in any order, but in which R_b doesn't apply, in any order, for any b distinct from a . Each model is therefore associated with a unique arrangement $Arr(M) = a$ where a is the arrangement such that the extension of R_a is non-empty in M . As with the Hilbert theory, whenever M and M' correspond to different arrangements — i.e. when $Arr(M) \neq Arr(M')$ — there is a sentence that is true in M but not M' , namely R_axyz where $a = Arr(M)$.

The principle we have appealed to in each case is the following principle. Suppose that \mathcal{L} is a theory with a class of intended models \mathcal{C} , that includes models representing each possible arrangement of the three particles (that is to say, the function Arr associating each model with an arrangement is surjective). Then if the primitives of \mathcal{L} express a physically complete set of fundamental relations, any two models representing different arrangements of the particles ought to be distinguishable by some sentence of \mathcal{L} :

DISTINGUISHABILITY: If M and M' correspond to different arrangements of the particles (i.e. $Arr(M) \neq Arr(M')$) then there is some closed sentence $\phi \in \mathcal{L}$ such that $M \models \phi$ and $M' \not\models \phi$.

This constraint is quite important and is effectively a way of saying, in model theoretic terms, that the arrangement of the particles supervenes on the facts expressible in \mathcal{L} . Without it one could not formulate an adequate physical theory: the forces particles exert on each other, for example, depend on the arrangement of those particles. In particular, particles arranged differently can behave differently. If the fundamental primitives do not distinguish between two possible arrangements of the particles then the behaviour of the particles will not be determined by kinds of facts we are taking to be fundamental. Neither the arrangements nor the motions of the particles will supervene on the distribution of the fundamental properties and relations; particles could be arranged differently in two worlds even when the two worlds agree about all the fundamental facts as stated in our fundamental language, \mathcal{L} .

We can now prove the following limitative theorem.

Theorem 2.1. *Let \mathcal{L} be a first or second order language, \mathcal{C} a collection of models of \mathcal{L} , and $Arr : \mathcal{C} \rightarrow \mathcal{A}$ a surjective association of arrangements to models. Suppose that our class of models also satisfies the distinguishability constraint. Then one of the following is true:*

1. \mathcal{L} has infinitely many non-logical primitives.
2. \mathcal{C} contains an infinite model.

¹⁹Proof: let the distances between x and y in a model M be denoted $d_M(x, y)$. If the arrangement of particles in M and M' differ then for some $a, b, c, d \in \{x, y, z\}$ the ratio $d_M(a, b)/d_M(c, d)$ and $d_{M'}(a, b)/d_{M'}(c, d)$ differ (if all these ratios agreed, they would be in the same arrangement). That means there are a pair of rational numbers, q and q' such that $q < d_M(a, b)/d_M(c, d) < q'$ but doesn't hold when M is substituted for M' . As we indicated earlier, it is quite easy to express facts about rational distance ratios using a single sentence.

This theorem is not particularly deep from a mathematical perspective.²⁰ However it does clarify the situation: the two most salient responses to the theorem are (i) to adopt infinitely many primitives, as with PRIMITIVISM, or (ii) to adopt an infinite ontology, as with the substantialist theory or the theory that postulates mathematical objects to account for the relations between the three particles. Note, however, that our theorem is quite general — we have allowed, for example, that which ‘extra’ objects exist can depend on the arrangement of the particles. Perhaps if the particles are nicely and symmetrically arranged one can characterize that arrangement with a finite set of primitives without postulating extra objects. What the theorem shows is that whatever your theory, if it has finitely many non-logical primitives, there will always be some arrangements for which one needs infinitely many objects present to distinguish those arrangement from other such arrangements.

Note that one could also question the constraints placed on the kind of language employed in our theorem: we allowed ourselves first order and second order resources, but one might think it possible to do better with modal resources. A particularly natural strategy would be to combine modal resources with second order resources. With a modest finitary principle of recombination for possible particles, guaranteeing that the outer domain of quantification is infinite, one could force the domain of the second order quantifiers to range over an infinite collection of intensions even whilst keeping the inner domain of the first order quantifiers finite at each world. This could in principle allow one to simulate the kinds things one would normally do with space-time or mathematical objects, whilst technically keeping our first-order ontology finite.²¹ Even if one grants that second order quantification does not in itself carry a commitment to abstracta (as I am inclined to think myself), such proposals are not without their own difficulties, although it is beyond the scope of this paper to survey them in full (but see the discussion in section 2 (on using properties to simulate space-time) and 9 (on modal approaches) of Field [12]. (See also Mundy [23] and Eddon [10] for property theoretic account of quantities, and Belot [6] for a fairly sophisticated example of a modal treatment of geometry.)²²

One could respond to these arguments not by introducing relations to space-time points or numbers, but by introducing some other infinite physical structure that, unlike space-

²⁰If \mathcal{L} has a finite signature, then for each finite cardinality n , there are only finitely many different isomorphism classes of models of \mathcal{L} of that cardinality (for an k place relation there are only 2^{n^k} possible relations over a domain of cardinality n , so there are $2^{n^{k_1} + \dots + n^{k_m}} \cdot n^l$ possible models in a language with l constants and predicates of arity k_1, \dots, k_m). It follows that if each model in \mathcal{C} is finite then there are at most countably many isomorphism classes of models of \mathcal{M} . But since there are uncountably many arrangements, Arr must map two models in the same isomorphism class to the same arrangement, contradicting DISTINGUISHABILITY.

²¹Here is an example of this kind of strategy, where we employ modal operators and higher order quantification into sentence position: one could adopt a primitive $\mathcal{A}(x, y, P)$ taking two terms and a sentence, roughly meaning P is an arrangement of x and y : in a possible worlds style model, for some distance d , P is the set of worlds at which the distance between x and y is d . From this one can define what it means for P to describe the arrangement of any finite collection of particles, and helping oneself to propositional quantification and a necessity operator, one can recover the congruence and betweenness relations. Thanks to Jeremy Goodman for discussion here.

²²According to the Mundy-Eddon view, there are infinitely many properties corresponding to each quantity (e.g. one property for each possible mass) and primitive higher-order relations between these properties that determine their quantitative structure. It is natural to view the Mundy-Eddon view as falling under the ‘infinite ontology’ branch of our dilemma, as particles are inheriting their quantitative structure by standing in relations to an infinite collection of mass properties. However, one could also formalize it in third-order logic with primitive third-order relations over second-order entities, in which case the theory is not within the remit of our theorem, and the first-order ontology could well be finite (even though the higher-order domains are infinite).

time, is invariant under Leibnizian symmetries. Although not relationism as traditionally conceived, one might hope to find an alternative to space-time from which distances and other quantities can be recovered without relinquishing NO SHIFTS.

There are, of course, lots of ways of going about this strategy, and the project is too broad to say anything too conclusive about its prospects. A particularly natural way to go about this is to take a group of *displacements* of the universe of material objects as a primitive physical entity in its own right. According to this theory, particles do not get locations in this space, but rather *pairs* of particles get assigned ‘locations’ — i.e. vectors in our space of parallel displacements — representing the displacement between them.²³ Crucially, this ‘location’ remains the same under Leibniz shifts.²⁴

Note however that these projects all involve accepting a form of substantivalism, even if not the familiar sort. Rather than evaluating the prospects for theories like this — a worthwhile project for another time — I want to focus on the possibility of carrying out this sort of idea within the confines of orthodox substantivalism.

3 Multiple Location

These sorts of considerations strike me as a good *prima facie* reason to be a substantivalist: by appealing to relations between material objects and space-time points we can recover enough structure to represent the distances and other quantitative properties and relations, without placing an undue significance on any particular mathematical representation of these facts. A relationist, on the other hand, must either take the role of abstracta in physics more seriously than seems wise or attempt to emulate substantivalism using certain kinds of higher order and modal resources.

For someone following this line of reasoning, it would not be unreasonable to take these considerations to also reflect negatively on the NO SHIFTS principle. After all, substantivalism is often taken to entail that there are differences between shifted worlds. Note, however, that the role substantivalism plays in fixing the geometrical structure of material objects is quite different from the role that regions of space-time play in the shift argument, which made some specific assumptions about the location relation.

To address the issue of geometric structure one postulates a manifold with its own intrinsic geometric structure. Geometric relations between material objects then exist only in a derivative sense: in virtue of material objects standing in some relation — call this the ‘location relation’ — to space-time regions that have the geometric properties in question. Nothing we have said so far, however, rules out the possibility that the location-like relations material objects bear to space-time regions are invariant under Leibnizian transformations.

²³Note that the order of the pair matters: if v is the displacement between p and q , $-v$ is the displacement between q and p .

²⁴For simplicity we might focus on a physical structure that is invariant under translations in space — the group of *parallel displacements* — however mathematically it straightforward to extend this idea to rotations using the notion of an *angular displacement*, and boosts and other continuous symmetries using other similar notions. A parallel displacement is just a vector quantity that points from point of space-time to another, telling us the displacement between the two points; this quantity is invariant under translations in the sense that if one pair of particles is a translation of another the displacement vector between the first pair is the same as the second pair. The group of parallel displacements of Galilean space-time is itself a four dimensional manifold, just like Galilean space-time, but it has more structure: it is a normed vector space, and thus unlike Galilean space, has a special point that is distinguished from the others (the 0 vector). The manifold is rich enough that the mathematical structure can be recovered in a Hilbert-Tarski-Field style setting.

Note that as I have introduced it above, the location relation is defined implicitly by its role as *that relation between objects and space-time from which the geometrical properties of objects can be recovered*. As such, one shouldn't assume that it corresponds to the pretheoretic notion of location we are used to employing when we are not doing fundamental metaphysics. The ordinary notion of location usually relates us to places — like Paris, the moon, and so on — and not regions of space-time.²⁵ With this caveat in mind, we can ask what it would mean for the world to remain unchanged by a uniform shift of the locations of each material object. Ignoring time for a moment, and restricting ourselves to a three-dimensional universe it would require the following:

SHIFT INVARIANCE: If x is located at y and y' is a Euclidean transform of y (a combination of shifts, rotations and reflections) then x is located at y' .

Later we will see that SHIFT INVARIANCE can be formulated very simply without appeal to Euclidean transforms (see section 3.4).

SHIFT INVARIANCE arguably follows from a certain conception of the location relation. As mentioned above, we are introducing the location relation by the job it plays, regardless of the distance from our pretheoretic notion. Assuming that an object's geometric properties are entirely determined by its shape, the job to be satisfied can be summarized by the idea that the shape of an object is given by the shape of its location.²⁶ Thus the location relation must satisfy the constraint that if it holds between a material object O and a region R , then O has the shape of R . The simplest relation satisfying this constraint is the relation ' O has the same shape as R ', and on this interpretation SHIFT INVARIANCE is true.²⁷ (Of course, there are also one-one relations that satisfy the job, relating each object to a unique region: but there are infinitely many of them that are equally good, and if the job description is our only criteria for choosing, to pick one of them seems arbitrary.)

Of course SHIFT INVARIANCE requires material objects to be multiply located in a fairly radical way. The crucial point is that if we are at a world in which objects are multiply located in this sort of way, then uniformly subjecting every object to a Euclidean transformation will leave each object with the locations it had before the transformation. Thus we do get to maintain the principle NO SHIFTS.

This idea generalises to other kinds of transformations. For example, if you thought that embiggenings don't generate genuine differences then perhaps a world in which the displacement between some particular point, p , and the location of each particle had been doubled (an embiggening around p) also results in a state of affairs no different from the one you started out with. (In this case, however, it is less clear that the transformation is a symmetry of the underlying physics.²⁸) In the special theory of relativity, the relevant transformations are Lorentz transformations. And in the context of general relativity, the

²⁵More importantly, objects are usually thought to be located at no more than one place on this conception.

²⁶More generally, we can read off geometrical relations between more than one object from the shape of their fusion.

²⁷Note that although one can identify the location relation with the relation of having the same shape, that needn't be the order of reduction: in particular, it's consistent to assume, as we have been, that space-time has its geometric structure intrinsically and that the shape of a material is given by its location(s).

²⁸In a world in which the distances between two point particles of mass m has been doubled, they will accelerate towards each other at different rates, because the forces between them are inversely proportional to the square of the distance between them. (Note that talk of 'doubling the distance between two points' must be treated with some care. One can of course change between units in a way that doesn't change the physics, but if we keep the scale the same and move each particle so that the distance between pair of particles is doubled we won't in general keep the physics the same, unless we also change other properties like their masses to compensate.)

transformations that seem natural are diffeomorphisms. In this latter setting the metric structure of the manifold seems like a contingent and changeable feature of the world, and thus transformations which preserve the differential structure but not the metric properties seem like the natural transformations to use. (Although, perhaps surprisingly, this means that properties like shape are not independent features of an objects location in GR, and only emerge in the presence of a gravitational field). Thus although classical physics, with its Galilean symmetry group, has been our focus the fundamental idea can be generalised to other theories in natural ways.

SHIFT INVARIANCE is not enough: for all we've said, a single fusion of particles could be located at a table-shaped region of space-time and a chair-shaped region of space-time, so long as it is located at every region of space-time that has that table or chair shape. To rule this out we want to require that any two locations of an object are Euclidean transforms of one another:

SHIFT EQUIVALENCE: If x is located at y and x is located at y' then there's a Euclidean transformation taking y to y' .

SHIFT EQUIVALENCE is required if we are to carry out the project described as GEOMETRIC REDUCTIONISM: if the shape of a material object, for example, is determined by the shapes of its locations, but SHIFT EQUIVALENCE failed, objects simply wouldn't have well-defined shapes.²⁹ If we want to satisfy NO SHIFTS and GEOMETRIC REDUCTIONISM at once, we had therefore better accept SHIFT INVARIANCE and SHIFT EQUIVALENCE. (Note, on the other hand, that we could in principle attempt to take geometric properties as primitive properties of material objects and regions, and attempt to reduce the location relation to them: an object is located at a region if and only if they both have the same shape. This would, of course, would guarantee SHIFT INVARIANCE and SHIFT EQUIVALENCE, but would require us to be a geometric primitivist.³⁰)

Finally, it is natural to require that the locations of mereological simples be themselves mereological simples (space-time points):

SIMPLE LOCATIONS: If x is mereologically simple then its locations are too.

This rules out extended simples.

With these three principles in place we are in a position to see, at least in outline, that the present view is in as good a position as the relationist is. For the relationist the arrangement of a collection of particles is given by the distances between each pair of particles. SIMPLE LOCATIONS tells us that the locations of a mereological simple are themselves mereologically simple; they are space-time points. Given a natural principle governing the interaction of parthood and location, to be discussed in the section 3.3, it follows that the location of the fusion of two mereological atoms is the fusion of two space-time points. (For convenience we shall call the fusion of two mereological atoms a 'diatom'.) By SHIFT EQUIVALENCE we know that all the locations of a diatom must be congruent to one another, and thus all consist of a fusion of two space-time points that are the same distance from one another. Thus the distance between two mereological atoms is always uniquely determined from the locations of the diatom they fuse, so that all distance facts between point particles (and thus the geometric properties of their fusions) can be recovered from their relation to the space-time manifold.

²⁹Note that if we weakened Euclidean transformations to diffeomorphisms, then objects won't have well-defined shapes. This is perhaps not unexpected given the consequences of general relativity for the naive notion of shape.

³⁰Thanks to [XXX] for discussion here.

3.1 Locations in Ordinary Language

Now one might wonder how the radical proliferation of locations posited by this kind of view is consistent with our experiences. After all, if I am looking at the General Sherman I see a solitary tree. If I were looking at a bilocated tree — a single tree with two exact locations — one expects to see two tree-like shapes belonging to the same tree. This is at least how philosophers typically describe paradigm cases involving multiple-location. By this reasoning, if the tree were multiply located at each of its Euclidean transforms we should expect to see a tree smeared out over all of space (if that is even possible to visualise), but this is emphatically not what we see.

One crucial difference between the bilocated tree and the present case is that in the former the observer is not herself bilocated. Indeed, in the former case the distance between the observer and the tree is not obviously well-defined: there is the distance between the observers location and the tree's first location, and the distance between the observers location and the tree's second location. In such a case SHIFT EQUIVALENCE could fail, since the locations of the fusion of the observer and the tree might not be congruent to one another if the observer is not symmetrically positioned between the two locations of the tree (this result can be demonstrated more rigorously in the theory of part and location to be developed in the next section).

As we have just seen, however, in the present setting in which we have SHIFT EQUIVALENCE the distance between the observers eyes and the tree is completely well-defined: it simply falls out of a property that the locations of the eye-tree fusion share — roughly, being a disconnected eye-tree shaped region whose connected parts are separated by a certain distance. Of course it would be a tall order to give a complete theory of perception in more fundamental terms, but surely whatever the correct theory is, the position of the tree in our visual field will depend only on the relative distances between our eyes, the trees and other background objects and will not depend on which particular regions of space-time the tree is located at. (One way to convince oneself that there is no conflict with experience, perhaps, is to appeal to the fact discussed in the last section. The pattern of locations on this picture determines all the facts about relative distances that the relationist needs. Thus the present view has the means to account for perceived relative distances if the relationist does.)

Similar puzzles can be warded off in analogous ways. One might think that it is possible to uniquely specify the General Sherman's location simply by gesturing towards a particular location. I could point at the General Sherman and say something like 'the General Sherman is located at that region of space-time over there, and not anywhere else'. But on the present view I haven't really succeeded in singling out a unique region of space-time, since my hand is multiply located the gesture I made is simultaneously related to every region of space-time with that shape.

All of this admittedly sounds a little wild at first. To put it in perspective, it might be worth recalling our opening remarks in which we distinguished between two kinds of motivations for substantivalism. The reasons I have been discussing so far have been rather theoretical — the simplest way to formulate a physical theory of distances and other physical quantities without giving particular mathematical objects undue physical significance appeals to space-time points. However one might have much more direct motivations for being a substantivalist: perhaps you think that regions of space-time play a more explicit role in our lives than I have been acknowledging. Perhaps when I wonder where my keys are I am implicitly wondering which region my keys occupy, for example, or when I learn where

Jones went on holiday I learn something about a particular region. On this picture facts about particular regions of space-time can be revealed simply by observing the objects that occupy those locations, and the theory that best explains our observations is the orthodox theory in which every object has no more than one location.

I suspect that locations in the everyday sense exist and are indeed the subject of our ordinary talk of ‘locations’. But I also suspect that locations in this sense aren’t fundamental entities. Perhaps they are places: ‘my pocket’ and ‘Paris’ both refer to places, and seem like reasonable answers to the question of where my keys are and where Jones went on holiday, respectively. Countries and pockets, like other material objects, will be multiply located in the more fundamental sense. Or perhaps places in the colloquial sense are material objects, but are ontologically ‘lightweight’ objects like holes and shadows — certain kinds of non-fundamental entities whose existence supervenes on the properties of other more fundamental objects. (Holes and shadows, then, are also multiply located in the fundamental sense.)

Relationists are often perfectly happy to engage in talk of places and locations in the way understood above. In this way the kind of substantivalist I have been describing can also accept this sort of non-fundamental talk of locations. But like the relationist, I deny that this way of talking is a perspicuous description of the fundamental metaphysics. Indeed, even the ordinary substantivalist should think twice before attempting to identify ordinary locations, like my pocket and Paris, with regions of space-time, since the former have much more interesting modal profiles than the latter.³¹

To a substantivalist motivated by the more direct sort of reasons mentioned above the view that people are multiply located in this radical manner might seem particularly bad. But the fact is that the methodology of taking our ordinary use of language at face value is a notoriously bad way to do *fundamental* metaphysics. One can see the present view as what one gets from taking a relationist picture and then expanding the ontology with a richly structured physical entity that plays the role in a theory of quantities that a mathematical object would otherwise have played. Counterintuitive results arise when one attempts to identify that physical entity with the non-fundamental way of talking about places and locations; but I think this kind of identification is ill-advised in the first place.

3.2 The combined theory of part and location

According to orthodox substantivalism there is a fairly simple way to determine the location of a complex object. If you know the locations of its parts — in particular, if you know the locations of all its *atomic* parts — then the location of the whole is just the fusion of those parts. In short: the locations of the atoms determine the location of the whole. This is consistent with the broader Humean thesis that the properties of and relations between point sized objects determine the properties of and relations between all objects. The fact alluded to, however, does not come for free — it is delivered by a plausible theory governing the interaction of parthood and location. One can concoct formal models in which the location of a whole is not determined by its atomic parts and which thus violate this theory. A pair of point particles, both located at space-time points, has a fusion that is located at the fusion of those two points. However it is simple enough to construct formal examples where the fusion is located elsewhere: on one not so far-fetched model the fusion might be located at a larger extended region containing the original two points — on this picture ordinary

³¹Of course some people do make these identifications — see footnote 2 — but they are not completely pain free (see e.g. Sider [27] §4.8).

objects like tables and chairs can be located at extended regions (with non-zero volume) even if they are finite fusions of zero-volume point particles.³² But there are models that deny any connection between the locations of the atomic parts and the fusion: it's logically consistent that the location of the fusion of two particles be any region whatsoever, whether containing the locations of the individual particles or not. However insofar as these kinds of situations are deemed pathological they should be ruled out by our theory of part and location.

The situation for the multiple location theory is not so simple. It seems clear that whatever kind of theory one adopts connecting parthood and location it cannot be true that the locations of a complex object are determined by the locations of its atomic parts. To see this, note that by SHIFT INVARIANCE, SHIFT EQUIVALENCE, and SIMPLE LOCATIONS, a mereological atom has an exact location at each space-time point and is located only at space-time points. It follows that any two mereological atoms have exactly the same locations. Yet a complex object can have much more interesting locations. Consider two different objects each composed of three atomic parts: one could have locations that are all equilateral triangles, whereas the other might have locations that are all the shape of some particular irregular triangle. These two composite objects have different locations, yet as we have seen their atomic parts have exactly the same locations. In general, knowing the locations of the particles is not enough to determine the locations of the things they compose.

Indeed it might at first seem that on this picture one must completely relinquish the principle that the location of a complex object is determined from the locations of its smaller parts. As it turns out things aren't this bad: the location of a whole is determined by the locations of its diatomic parts (objects with exactly two proper parts) — but in order to see this we must develop the theory of parthood and location a little more.³³

To that end let us begin with the *pure* theory of location and part: the theory one gets by looking only at the relation between the parthood and location relations, ignoring any geometrical constraints such as SHIFT EQUIVALENCE and SHIFT INVARIANCE. Since, to my knowledge, no-one has given a thorough analysis of the relation between locations and parts among multiply located objects we shall need to develop a little bit of theory.

A natural constraint to impose on our theory is that it should have what I'll call the *classical theory of part and location* as a special case: it should entail that theory given the assumption that every object has a unique location. The classical theory effectively says that the location relation determines an isomorphism between the parts of a thing and the parts of its location (this theory is sometimes called 'mereological harmony'; see Uzquiano [31] and Saucedo [26]). For simplicity I shall assume the classical extensional mereology. I shall write $x + y$ to denote the fusion of x and y and $x \setminus y$ to denote the relative complement of y from x , which exists whenever x is not a part of y (formally, it is the fusion of x 's parts disjoint from y).

1. The domain of objects that have locations is closed under fusions and parthood

³²The fusion of the locations of the particles composing a table is an extremely disconnected object — the gaps between the particles is significantly greater than the sizes of the particles — so this intuition holds even if we do not assume that particles are point sized. On this view the table itself is solid, connected object merely having a location that contains the fusion of locations of the particles. See Fine [13].

³³In addition to the above, there are also quite general problems surrounding the interaction of mereology and location when multiple location is permitted; see Kleinschmidt [16]. The following is an attempt to formulate a simple logic of mereology and location that is not subject to these sorts of worries.

2. The location relation, L , is functional (on the domain of objects that have locations). We will write $l(x)$, to denote the unique y such that Lxy .
3. l preserves atoms: if a is mereologically atomic then $l(a)$ is mereologically atomic.
4. l preserves relative complements $l(x \setminus y) = l(x) \setminus l(y)$.
5. l preserves fusions: for any x and y , $l(x+y) = l(x) + l(y)$. (More generally, the location of the fusion of some things is the fusion of the locations of those things.)³⁴

Regarding the theory of parthood as it applies to both material objects and regions, the classical theory consists of classical extensional mereology. On the intended interpretation both regions and material objects individually and jointly form a complete Boolean algebra with the bottom element removed.

The classical theory rules out many of the interesting possibilities that have traditionally preoccupied philosophers interested in location. (5), for example, says that an object is located at the sum of its parts locations; thus ruling out the pathological example we opened with. (1) ensures that fusions and parts of located objects have locations, (2) rules out the possibility of multiple location, (3) rules out extended simples and (4) effectively rules out colocation and partial colocation (when mereologically disjoint objects are located at overlapping regions), and (3), (4) and (5) together rule out unextended complexes.

These principles jointly entail that the location function preserves other important mereological properties: for example that l preserves parthood and disjointness, and that no two things can have the same location.³⁵

Although many of the possibilities suggested above are worthy of philosophical attention, our purpose at present is to explore the consequences of multiple location, and the best way to do this is to screen off these other kinds of non-standard behaviour. Thus the theory we will consider is the minimal generalisation of the classical theory that allows for multiple location.

If l is a function that satisfies these conditions we shall call it a *location function*. The intended models of the classical theory of location and part are therefore models in which the location relation is a location function. (More formally the intended models of the classical theory of location and part consist of a tuple (D, \leq, o, l) where (D, \leq) is a standard model of classical extensional mereology, o is an element of D whose improper parts represent the objects that have locations, and l is a location function whose domain consists of the improper parts of o and whose range is disjoint from o .) In order to develop the multiple location theory further, we need to start by specifying the intended models of that theory.

By way of motivating our choice of model, let us start with a simple example. Consider a simple lego construction consisting of three lego bricks, B1-B3, that have been put together in a pyramid shape, with two at the bottom and one brick on top holding them together. Now suppose further that this lego pyramid is bilocated. By examining our intuitions about

³⁴The parenthetical part of (5) is stated using plural quantification. If a first order axiomatisation of this theory is sought, one should replace this axiom with a schema saying: if x fuses the ϕ s then x s location is the fusion of the locations of the ϕ s. Note that without the parenthetical these principles do not entail that l preserves arbitrary fusions.

³⁵(5) entails that l preserves parthood: if x is a part of y then $x+y = y$ and thus $l(x) + l(y) = l(y)$ by (5), which means that $l(x)$ is a part of $l(y)$. Without this principle one could be located in one city even when ones arms, legs, torso and head are located in another. By contrast (4) ensures that l preserves disjointness: if x is disjoint for y then $x \setminus y = x$ and so $l(x) \setminus l(y) = l(x)$ which means that $l(x)$ is disjoint from $l(y)$. It also ensures that no two things can have the same location, for if x is distinct from y then either $x \setminus y$ or $y \setminus x$ exists, but if $l(x) = l(y)$ then neither $l(x) \setminus l(y)$ nor $l(y) \setminus l(x)$ would exist.

this simple case we shall attempt to tease out some general principles connecting the location and parthood relations.

To keep our intuitions clean we shall make a few simplifying assumptions. (i) The three bricks themselves might be composed of smaller parts — indeed assuming they are not extended simples they must — but for simplicity we shall pretend that these three bricks and their fusions are the only parts of the pyramid. (ii) We shall assume that the two locations of the lego pyramid are both congruent to one another (in accordance with SHIFT EQUIVALENCE). This geometric constraint is not required by the pure theory of part and location — the different locations of the pyramid could in principle have been any other shape or size (although it would then become unclear what grounds we have for calling it a lego pyramid, if some of its locations did not have that shape). (iii) We shall assume that the bricks B1-B3 that compose the pyramid also have two locations each, and that these two locations are congruent. (iv) We shall also assume the three bricks are arranged symmetrically at each location — one could in principle have had B1 be the ‘top’ brick relative to one location and B2 on top at the other. Again this is completely consistent with the pure theory of part and location we are exploring in this section. (v) Lastly we shall assume for simplicity that the two locations of the whole pyramid do not overlap — to make things vivid we’ll suppose one of the pyramid’s locations is entirely contained within a box, and the other within a jar.³⁶ Our theory ultimately won’t require this, and of course any constraint like this will have to be dropped if we are to accommodate SHIFT INVARIANCE.

The interaction between parthood and location here is relatively clear: there are 7 (i.e. $2^3 - 1$) parts of the pyramid: B1, B2, B3, B1+B2, B2+B3, B1+B3 and B1+B2+B3. Each of these parts, including the pyramid itself, has two locations — a box location and a jar location. Insofar as we are generalising the classical theory, then the box location of B1+B2 should just be the fusion of the box location of B1 and the box location of B2, and similarly for the other three composite parts of the pyramid. Since we have only finitely many objects this in effect gets us the classical theory of location and part when we restrict quantification over locations to locations within the box. None of the deviant behaviour ruled out by (1)-(5) ought to arise when restricting ourselves to box locations (ignoring, for the moment, that the bricks are extended simples): thus there ought to be a location function — a function satisfying (1)-(5) — mapping parts of the pyramid to regions within the box. Parallel reasoning suggests that there should be a location function mapping the pyramid and its parts to regions inside the jar. These location functions, of course, are also defined on material objects disjoint from the pyramid; however assuming there is no other multiple-location going on, we can assume that the two location functions agree about the locations of every material object disjoint from the pyramid. An arbitrary object is thus located at a region iff one (or both) of these two location functions maps the object to that region: formally this means that the location relation is just the union of these two location functions (recalling that both functions and relations are sets of ordered pairs).

This example is hopefully instructive enough as to make the following model of a location relation seem particularly natural


A proper location relation, L , is a union of location functions, where each location function is a mapping of material objects to regions of space-time satisfying (1)-(5).

³⁶Note if the two locations of the pyramid did overlap, that would be completely consistent with our ban on colocation. That ban rules out two objects being entirely located at the same region, or entirely located at overlapping regions: but this is a case where only *one* object is entirely located at two overlapping regions, and so is not subject to this ban.


Thus any location relation, L , can be decomposed into a set of location functions roughly telling us how objects and their parts are located relative to certain locations. The intended models therefore consist of all tuples (D, \leq, o, L) where (D, \leq) is a standard model of classical extensional mereology, o is an element of D whose improper parts represent the objects that have locations, and L is a proper location relation whose domain consists of the improper parts of o and whose range is disjoint from o .

It should be stressed that the decomposition of L into location functions is not unique: one can have two sets of location functions that have the same union. Here is a very simple example of this phenomenon. Suppose that two point particles a and b are both triply located at each of three space-time points, x, y, z , arranged in an equilateral triangle. Suppose further that the fusion $a+b$ is multiply located at the three locations $x+y, y+z, x+z$ respectively (note that this stipulation is consistent with the constraint that L be a union of location functions).

This situation can be represented by three location functions as follows: Here each box



represents how a and b are located in space relative to each location function. Note that it can also be represented by the following three location functions:



The unions of these two sets of location functions are identical, and locate a and b in exactly the way we described above.³⁷ It should also be obvious that the union of all six location functions gives us yet another representation of the same relation. However this latter representation is special in the sense that it is the largest set of location functions whose union is L . Thus although there is no unique decomposition of a location relation into location functions, there is always a unique maximal decomposition: the set of location functions that are subsets of the location relation. When we talk about the decomposition of a location relation we shall always mean the maximal decomposition.

This observation suggests that while location functions are a useful way to specify the structure of the location relation they don't have any independent reality; indeed in the next sections we shall look at ways of characterising the location relation directly without appealing to location functions. It is worth comparing our theory with *prima facie* similar responses to the shift argument that place more importance on the role of location functions. For example, Jeff Russell [25] has recently defended the view that material objects have unique locations but it is a completely indeterminate or non-factual matter which exact locations they have. Each 'precisification' of the location relation, on this view, is a location

³⁷We also constructed it so that it satisfied SHIFT EQUIVALENCE.

function, telling us where each object is located according to that precisification. It is interesting to note that if one were to apply this model to the above example there would be a genuine difference between a world where it is indeterminate whether the locations of a, b and $a + b$ are given by the first, second or third location functions, and another world where it is indeterminate whether the locations of a, b and $a + b$ are given by the fourth, fifth or sixth (so that the admissible precisifications of ‘located at’ are given by the first three locations functions in the first world, and the second three in the second world). On the picture where they’re multiply located, rather than indeterminately located, there would be no difference between these two possibilities because they determine the same location relations.

Ultimately these distinctions get washed out in Russell’s theory: he accepts a principle analogous to SHIFT INVARIANCE that effectively forces us to choose the maximal set of precisifications. But this demonstrates that our theory of multiple location is at least conceptually very different; it does not make the kinds of invisible distinctions that would have to be acknowledged if we started with a set of admissible location functions instead of the single location relation that is their union.

3.3 The Theory of Part and Location

Now that we have specified the intended models of the theory, let us turn to axiomatising it. A major choice-point is whether to pursue a first or second order axiomatisation. In what follows I shall present both, however there are several substantial technical questions concerning the first order theories that remain open. I outline these in further detail in the appendix. These mainly concern completeness (is every sentence true in all intended models provable in our system), and the status of representation theorems (can every relation that satisfies the axioms of the theory be expressed as a union of location functions).

The plural language augments first order logic with plural variables, written xx, yy, zz , and so on, a plural quantifier $\forall xx$ and a singular-plural relation $x \prec xx$ meaning that ‘ x is one of the xx ’. In both the first and second-order cases, the non-logical vocabulary consists of a binary relation, L , a unary predicate O , and a binary function symbol $+$. L represents the location relation, O applies to material objects, and $+$ is a fusion operation. We assume, for simplicity, that every object is either a material object or a region of space-time so that regions of space-time can be defined as $\neg O x$.

We have chosen to take the binary fusion operator as primitive for convenience, but note that many of the more familiar mereological relations can easily be defined from it. Parthood is defined by $x \leq y =_{df} x + y = y$, overlap by $x \circ y =_{df} \exists z(z \leq x \wedge z \leq y)$, and atomhood $At(x) =_{df} \forall y(y \leq x \rightarrow x = y)$. We write $xx \subseteq yy$ a short for $\forall z(z \prec xx \rightarrow z \prec yy)$. Finally, if x and y overlap, it is also convenient to have a term $x \sqcap y$ — the *product* of x and y — that denotes the fusion of things that are parts of both x and y , and a term $Fus(xx)$ which, whenever there is at least one xx , denotes the fusion of the xx .

We may now start laying out the axioms of this theory. We shall start with the plural version, which results from adding to a standard axiomatisation of plural logic, the following (self-explanatory) principles:

The axioms of classical mereology (in terms of $+$).³⁸

³⁸From the first three axioms and our definition of parthood one can prove the reflexivity, transitivity and anti-symmetry of the parthood relation. (For example, to show anti-symmetry suppose that $x \leq y$ and $y \leq x$. Thus by definition $x + y = y$ and $y + x = x$. But by the second axiom $x + y = y + x$ so $x = y$. The

$$\begin{aligned}
x + x &= x \\
x + y &= y + x \\
x + (y + z) &= (x + y) + z \\
\forall x \exists x \forall y (x \circ y &\leftrightarrow \exists z (z \prec x \wedge y \circ z))
\end{aligned}$$

All and only material objects are located somewhere.

$$Ox \leftrightarrow \exists y Lxy$$

Locations aren't material objects

$$Lxy \rightarrow \neg Oy$$

Atoms are located at space time points

$$At(x) \wedge Lxy \rightarrow At(y).$$

In the classical theory of part and location, we had an axiom that said that the location relation preserved fusions: $l(x + y) = l(x) + l(y)$. Our approach to axiomatising the theory of part and location will be to produce a suitable analogue of this axiom. However one must be careful: it is not true, for example, that the locations of $x + y$ are all possible fusions of the locations of x and the locations of y . In the case of the lego pyramid, for example, the fusion of B1's box location and B2's jar location is neither wholly located in the box or the jar, and is thus not a location of B1+B2.

Here is something that is true, however. If the pyramid (the fusion $B1 + B2 + B3$) is a located at a region R , then there is a way of cutting R up into three disjoint pieces $R1, R2$ and $R3$ such that (i) $B1$ is located at $R1$, $B2$ at $R2$ and $B3$ at $R3$, and (ii) $B1 + B2$ is located at $R1 + R2$, $B1 + B3$ at $R1 + R3$, $B2 + B3$ at $R2 + R3$... and (iii) $B1 + B2 + B3$ is located at $R1 + R2 + R3$.

This idea can be generalised in the following way. Suppose that an arbitrary object x is located at a region y . Then for any way of cutting x up into smaller disjoint pieces, $(x_i)_{i \in I}$ (indexed by some set I), there is a way of cutting up y into an equal number of little pieces $(y_i)_{i \in I}$ such that whenever $J \subset I$, $Fus(\{x_j \mid j \in J\})$ is located at $Fus(\{y_j \mid j \in J\})$. Thus by letting J range over singletons, we can see that each part of x , x_i say, is located at the corresponding part of y , y_i . Moreover, each fusion of these parts of x (say, $x_i + x_j + x_k$) are located at the fusion of the corresponding parts of y , (that is, $y_i + y_j + y_k$). (And similarly for infinite collections of these parts.)

This brings us to the final and most important axiom of our theory, which effectively encodes this thought in plural logic. A partition of an object, x , are some things xx which are pairwise disjoint ($\forall xy(x \prec xx \wedge y \prec xx) \rightarrow (x \circ y \rightarrow x = y)$) but fuse to x ($Fus(xx) = x$). A partition xx divides x if and only if x is a fusion of some of the xx . Finally write $mat(x)$ for the material part of x (i.e. $x \sqcap Fus(\langle x : Ox \rangle)$) and $reg(x)$ for the space-time part of x (i.e. $x \sqcap Fus(\langle x : \neg Ox \rangle)$).³⁹ The final axioms says:

other principles are similarly straightforward.) Thus every principle of classical mereology is provable given our definition of parthood (the fusion principle needs no modification). Conversely, the below principles can be proven in classical mereology if we add the linking principle that $x + y = z$ iff z is the fusion of x and y : $x + y = z \leftrightarrow \forall u(u \circ z \leftrightarrow (u \circ x \vee u \circ y))$.

³⁹Here $\langle x : \phi x \rangle$ is a plural expression for the ϕs .

ARBITRARY PARTITIONING Suppose that x is located at y and that the xx are a partition of the material objects that divides x . Then there are some things, the zz , that are a partition and that divides $x + y$, such that:

- (a) If z and w are one of the zz , and $mat(z) = mat(w)$ or $reg(z) = reg(w)$ then $z = w$.
- (b) If the $ww \subseteq zz$ then $mat(Fus(w))$ is located at $reg(Fus(w))$.
- (c) If $z \prec zz$ and $mat(z)$ is atomic, $reg(z)$ is atomic.

Our formalisation of the intuitive thought, which uses indexing sets, is a little less direct, since we are restricted to a plural language. The zz effectively encode a pairing between the given partition of x and some partition of y , and a little reflection should reveal that ARBITRARY PARTITIONING correctly captures the thought explained above.

The above plural axiomatisation has an important property: it accurately captures the intended models of the pure theory of part with multiple-location. Any full model of the axioms is one in which L can be represented as a union of location functions.⁴⁰

In order to state a first-order version of this theory we must first replace the mereological composition axiom above with a schema:

$$\exists z\phi \rightarrow \exists x\forall y(x \circ y \leftrightarrow \exists z(\phi \wedge y \circ z))$$

We must also find a suitable first-order replacement for ARBITRARY PARTITIONING. The most straightforward substitute is the following finitary version:

PARTITIONING: If x is located at y , and x_1, \dots, x_n partition x , then there are y_1, \dots, y_n such that that each x_i is located at y_i , each $x_i + x_j$ is located at $y_i + y_j$, each $x_i + x_j + x_k$ is located at $y_i + y_j + y_k$ (and so on).

Unlike ARBITRARY PARTITIONING, PARTITIONING is a schema, with a different instance for each choice of n . That this schema really is expressible in first-order logic is shown in the appendix.

Here the logical issues are not as clean as with the second-order case. In any model of the theory in which the extension of O is finite, the extension of L is a union of location functions. However there are unintended models when the extension of O is infinite. Two natural questions then present themselves: are there *any* first order axiomatisations that rule out unintended models, and are there any first-order axiomatisations that are sound and complete for the class of intended models (even if such axiomatisations admit unintended models)? Both these questions are explored further in the appendix.

⁴⁰A model of a plural language is *full* if the plural quantifiers range over every subset of the domain. We show that L is the union of the location functions that are subsets of L (here we use L to denote the interpretation of the location function in a model). In particular this means showing that if Lxy then some location function that is a subset of L maps x to y (we show this under the assumption of atomism; a more intricate argument is needed without the assumption). Suppose that x is located at y and consider the partition that consists of all the material atoms. Arbitrary partitioning guarantees that there is a bijection f between these atoms and the atoms of a region of space-time containing y , whose graph is determined by material and space-time parts of the elements of the partition zz . This bijection will be such that (i) the fusion of any set of those atoms, X , is located at the fusion of the image of X under f , $f(X)$. If we let $l(a)$ be the fusion of the image of a atoms under f then l is a location function that is a subset of L . Moreover, since zz divides $x + y$, l maps x to y as required.

3.4 Combining the Theory with Shift Invariance

We can eliminate reference to the mathematical notion of a Euclidean transformation in our statements of SHIFT INVARIANCE and SHIFT EQUIVALENCE, so that our entire theory can be stated in the internal language of congruence, betweenness, location and part.

In fact the statement of both principles takes a particularly simple form, and turns out to be a restriction of both principles to diatoms. For example, ‘diatomic shift invariance’ says that if a diatom x is located at $p + q$ (a fusion of two space-time points) then it is located at every Euclidean transform of $p + q$. Note however that $p' + q'$ is a Euclidean transform of $p + q$ iff p and q are congruent to p' and q' , which is something we can state in our fundamental vocabulary consisting of congruence and betweenness relations. Diatomic shift equivalence similarly becomes the claim that if x is located at $p + q$ and $p' + q'$ then p and q are congruent to p' and q' . Putting both together we get

CONGRUENT DIATOMS: $Lx(p + q) \rightarrow (Lx(p' + q') \leftrightarrow Cpq p'q')$ where p, q, p' and q' range over space-time points.

To see that this principle is adequate, we need to know that if the diatoms are located at all and only the Euclidean shifts of their locations then every object is located at all and only the Euclidean shifts of their locations. In other words, we need to know that CONGRUENT DIATOMS entails SHIFT INVARIANCE and SHIFT EQUIVALENCE. In fact this follows from a fact we mentioned earlier and which we are now in a position to prove: the locations of a whole are determined by the locations of its diatomic parts (even though, as we saw earlier, they are not determined by the locations of their atomic parts).

Theorem 3.1 (The Locations of the Diatoms Determines the Locations of Everything). *Let L be a proper location function (a union of location functions), between some domain of material objects and Euclidean space E^3 . Suppose that for every diatom there is a pair of space-time points such that the diatom’s locations are given by the collection of all pairs of space-time points congruent to that pair. Then for any material object, there is some region such that the object’s locations are given by the collection of all regions congruent to that region..*

Proof. Suppose that that L is a union of a set of locations functions, F , and that d is some metric that coheres with the congruence and betweenness structure on E^3 . We want to show that if Lxy and Lxz then y is congruent to z : there is some Euclidean transformation that maps y to z . It is a standard result that y and z are Euclidean transformations of one another if and only if there is a distance preserving bijection (an ‘isometry’) between the points in y and the points in z .

From the fact that Lxy and Lxz it follows that there are location functions $f, g \in F$ such that $f(x) = y$ and $g(x) = z$. I claim that the mapping $\iota = g \circ f^{-1}$ restricted to points is an isometry between points in y to points in z . One can see by inspection that it maps y to z . Suppose that p and q are two points in y and that $\iota(p) = p'$ and $\iota(q) = q'$. Now let $a = f^{-1}(p)$ and $b = f^{-1}(q)$. It follows that $g(a) = p'$ and $g(b) = q'$. Note also that a and b must be mereological atoms because they are mapped to atoms in E^3 by a location function. Thus $a + b$ is a diatom, and it is located at $p + q$ by f (note that $f(a + b) = f(a) + f(b) = p + q$) and located at $p' + q'$ by g (since $g(a + b) = g(a) + g(b) = p' + q'$). By CONGRUENT DIATOMS it follows that $p + q$ is congruent to $p' + q'$, which means that $d(p, q) = d(p', q') = d(\iota(p), \iota(q))$ as required of an isometry. \square

A small modification to the argument must be made to extend to the four-dimensional case: regions are transforms of each other in the relevant sense iff there is an isometry that also preserves simultaneity and non-simultaneity of each pair of points.

4 Conclusion

We have considered two sorts of theories of the geometry of material objects: theories in which they inherit their structure by their relation to a space-time manifold which has its geometric structure internally, and theories in which they inherit their structure by their relation to a platonic realm of abstracta. We have seen that the most straightforward versions of both these views give rise to certain kinds of objectionable invisible structure. That does not mean we should reject these approaches, for they may have other theoretical virtues that could not be achieved otherwise. However I have outlined a theory that can provide an adequate account of the geometry of physical objects that does without these invisible differences, and is otherwise quite simple. To be sure, the view does not describe (nor does it purport to describe) our untutored ideas about the nature of space-time as it relates to our ordinary talk of locations. But it strikes me that this talk is just a convenient shorthand for describing facts about the *relative* locations of different objects (much like the physicists use of coordinate systems) and should not be taken seriously as a guide to how things are fundamentally.

5 Appendix A: Fields

Fields, as they are ordinarily conceived, are properties of space-time: a field determines a field value for each point in space-time, which specifies the strength (and possibly also direction) of the field at that point. On this view the strength of the field values at a particle's location determines the behaviour of that particle. This picture, however, is fundamentally at odds with the present view, since each particle has many locations, and so no unique field value can be associated with a particle. This obstacle is not fatal to the present approach: Newtonian mechanics can be formulated as an 'action at a distance' theory — a theory that works directly with the forces acting on each particle by other particles without invoking fields to mediate these forces. However this approach is not in the spirit of (if not incompatible with) the idea that forces and other quantities usually represented by mathematical objects are fundamentally reducible to more basic relations between concrete entities like space-time points. For example, the action at a distance theory assigns a number to each ordered pair of particles telling us the force one exerts on the other. A problem similar to that posed for relationism then arises: unless there are a continuum of particles between each pair of particles, relations encoding the forces between the particles alone will be insufficient to specify all the possible kinds of forces a particle can exert on another. If one could consider the forces exerted at a continuum of points between the two particles (as one might in some sort of field theory) we could employ the methods discussed in section 2.2 to eliminate numbers. At any rate, the use of space-time fields are so pervasive in modern physics it would be remiss not to treat them.

In what follows I'm going to focus on strategies, in the spirit of those proposed by Field [11], for demathematizing scalar and vector fields by reducing them to certain kinds of congruence and betweenness facts. This strategy comes in two parts, the first of which is to note that most vector fields of physical interest can be equivalently described by a scalar

field, so that our problem can in effect be reduced to the problem of giving an account of scalar fields. The second is to see that a scalar field can be represented, up to an affine transformation, by a pair of congruence and betweenness relations.

If a vector field is ‘conservative’ — a property most physical fields share — then it can be represented as the gradient of a scalar field.⁴¹ To get a feel for this result it helps to visualise the special case of a scalar field on two dimensional space: one can picture this as a kind of hilly terrain imposed over the surface where the height of the field above the surface represents the magnitude of the scalar field. The slope of the terrain at a point — the vector pointing downhill with a magnitude proportional to the steepness — will be a vector field of the requisite kind. This idea easily generalises to higher dimensions. When a vector field is generated by a scalar field in this way, we call the scalar field a ‘potential’ for the vector field.

It should be noted that many different scalar fields generate the same vector field — one can uniformly raise or lower the height of a terrain without changing the direction or magnitude of the slope at each point. For this reason people normally don’t consider the potential scalar field to be fundamental. When the laws of physics depend only the values of a vector field, and thus do not discriminate between different scalar fields that generate that vector field, it is extremely natural to think that the only physically real distinctions are those that give rise to different assignments of vectors to points. It thus seems *prima facie* inadvisable to attempt to reduce vector fields to scalar fields. However our strategy is to reduce both scalar and vector fields to congruence and betweenness relations, so, for example, the potential betweenness relation would say ‘the potential at x is between the potential at y and z ’. Interestingly, in this setting the correspondence between vector fields and the potential congruence and betweenness facts becomes one-to-one. In particular, one can change the height of a potential without either changing the vector field it generates or the congruence and betweenness facts it generates. Thus the strategy of reducing conservative vector fields to facts about potential congruence and betweenness at space-time points seems like a promising place to start.

However once we adopt SHIFT INVARIANCE we quickly encounter a completely independent difficulty that has nothing to do with the project of demathematizing fields. To illustrate the problem, consider a single particle with external forces acting on it. We would normally represent the gravitational potential by a distribution of quantitative properties over space-time points and determine the particle’s motion by the distribution of these properties in a neighborhood of the particle’s location. If the particle is multiply located at every space-time point this method breaks down: the gradient of the potential is different at each point of space, so we can’t tell anything about the force acting on a particle just by looking at its locations. It is natural to think that this means that fields have to be multiply located as well: if a point p has a certain gravitational potential, every Galilean transformation of p must have that potential as well. But this is not sufficient either, for it would entail that every field value that’s had anywhere is had by every space-time point at once (since every space-time point is a Galilean transform of every other space-time point). The locations of the field has to be somehow correlated with the locations of the material objects for this to work.⁴² The following is an exploration of one way of meeting this challenge.

⁴¹A vector field is conservative if the ‘work-done’ to get from one point to another does not depend on the path you take. Most fields that have a chance of representing real physical fields are conservative our restriction won’t cost us much in the way of generality. Note, though, that an arbitrary vector field on an n -dimensional manifold can be decomposed into n scalar fields representing its n -components (although this decomposition is highly non-unique). So if needed, a more general reduction procedure is available.

⁴²Let me note in passing that one way to do this would be to expand our ontology of material objects to

5.1 Demathematizing fields in orthodox Newtonian physics

We shall start by getting a clearer handle on the representation of forces in the ordinary substantialist setting, in which each particle has a single location. Let us begin by considering a very simplified example consisting of two point particles, a and b , constrained to one dimension with no external forces acting on them. Both a and b generate a force field defined over space: the value of the field generated by a at the location of b determines how quickly and in what direction b will accelerate. Such forces can be equivalently represented by a scalar potential (the potential energy): in this case it will be the values of the potential in a *neighborhood* of a and of b that will determine their respective motions.⁴³ Now we can imagine what would happen if we were to uniformly translate the positions of the two particles in some direction by some fixed amount without also shifting the fields. The result would not be legal: each force field has a source — a point from which the field emanates — and these sources would not match the locations of the particles. In particular the force acting on a and b respectively would depend on their distances from the sources, and not on the distance from each other as it ought to. In order to preserve the laws we must translate a , b and the gravitational field as well. Similarly, when the two particles move in a lawlike manner, we must also provide laws telling us how the field values change, so that the field sources ‘keep up’ with the particles.

A metaphysically distinct (although mathematically equivalent) way to represent our two particle example is to imagine the system having a single location in a higher-dimensional space called *configuration space*, where each point in configuration space represents all the particles locations in ordinary space. Recall that both particles in our example reside in a 1-dimensional space, and there are only two of them, so in this case configuration space is 2 dimensional. (More generally, when we are considering all three dimensions, and there are N particles, configuration space will have $3N$ dimensions.) The gravitational forces can similarly be represented by a vector field on this 2-dimensional space, with the first component of the vector at a point representing the force a exerts on b and the second component representing the force b exerts on a when a and b are located according to that point in configuration space. As before, this vector field can be equivalently represented by a scalar potential. A uniform shift of the locations of a and b in configuration space corresponds to moving the location of the system diagonally in configuration space (north east or south west, as it were). Unlike in the previous example, we do not also need to shift the field, since the value of the gravitational field on configuration space is constant along all diagonal lines of this sort. The forces between two particles depends only on the relative distances between the particles, so the field values on configuration space must be constant along all paths in configuration space that leave the relative distances between particles fixed.⁴⁴ (Although

include fields as well, consisting of point sized parts. Field values can then be construed not as properties of space-time points, but as properties of the point-like parts of the field. Thus the field values are multiply located in a way that correlates with locations of the field itself, and any other material objects we might want to include. This strategy is not without contention, however. Since each material object has a unique ‘location’ within a given field, this maneuver potentially reinstates the problems we originally had with space-time. For example, Arntzenius [2] points out (in the context of a discussion of relationism) that you can shift the field values at each point of a field to some other point on the field in a uniform fashion, in a way completely analogous to a Leibniz shift on space-time.

⁴³Sometimes it is natural to consider a single vector field defined over space — the gravitational field — telling us the force per unit mass that would be exerted at each point of space-time due to every other particle. I have chosen to illustrate things with two distinct fields, rather than merging them into one, because it makes for a more straightforward comparison with views consider later.

⁴⁴In more detail, the value of the force at (x, y) — the point where a is at x and b at y — is $(-Gm_a m_b)/(x -$

we are presently focusing on Newtonian mechanics in the ordinary setting, one can already see how this second way of thinking might be more friendly to the multi-locational view I am endorsing.)

These two examples suggest two possible ontologies one might adopt. According to the first ontology one has two particles whose fundamental properties include their positions and masses, and two vector force fields (or equivalently, scalar potentials) whose fundamental properties include its strength at each point in ordinary space. To specify the dynamical evolution of this system we need two kinds of laws. One of the laws will tell us how the positions of the particles change over time — these laws tell us how positions change in terms of the strength of the field at (or at a neighborhood) of those positions. Another kind of law is needed to specify how the fields change their field values over time — these laws tell us how the field values change in terms of the positions of the particles. Conceptually this seems a little unsatisfactory, since there is an obvious circularity, and it's a mathematically non-trivial fact that the circularity is not vicious. This picture also seems to be unnecessarily complex: we have to provide both laws telling us how the positions of the particles evolve (depending on the strength of the field) and laws telling us how the gravitational field evolves (depending on the positions of the particles). This might seem especially unsatisfactory if one thought that the existence of fields were only postulated to explain the motions of particles; on the present picture also needs to explain the motions of the fields, and to do so one invokes the motions of the particles.

According to the second ontology we have particles as before, except now the thing playing the role of the two force fields (or potentials) is a single field over configuration space. In this setting the force field on configuration space remains constant over time, and the only thing that changes is the systems position in configuration space. Thus in this setting we only need laws telling us how the position of the particles change in configuration space in terms of a constant field.

Let us develop the second idea a little more. In the general setting with N particles and three spatial dimensions, the field can be represented by a scalar potential field U on a $3N$ dimensional configuration space, and the system gets assigned a position in that space. To get an intuitive handle on U , suppose that a 's position is described by the i th argument of U , and imagine that somehow the positions of every particle except for a has been frozen in place in positions $p_1, \dots, p_{i-1}, p_{i+1}, \dots, p_N$. We can then imagine the forces that would be exerted on a if a was located at each point p : these forces can be represented by the potential at each point in space. This is a function of p , U_i , that is given exactly by our N place function with $N-1$ of its degrees of freedom removed by filling them with the parameters $p_1, \dots, p_{i-1}, p_{i+1}, \dots, p_N$ (that is, $U_i(p) = U(p_1, \dots, p_{i-1}, p, p_{i+1}, \dots, p_N)$).⁴⁵ We shall simplify this idea in three stages: we shall firstly show how to do without configuration space substantivalism, then we'll show how to rid our primitives of haecceitistic involvement of particular particles, and finally we'll show how to understand these fields without fundamentally involving numbers in our description.

The first thing to note about this picture is that it seems to commit us to configuration space substantivalism: the view that configuration space is just as real and concrete as ordinary space, or perhaps that it should replace ordinary space. This is not in itself a problem, and some have suggested configuration space substantivalism for independent

$y)^2, Gm_b m_a / (x - y)^2$). It is the same value at $(x + c, y + c)$ since $((x + c) - (y + c))^2 = (x - y)^2$.

⁴⁵Note, of course, that to determine the motion of a , one only needs to know the field values in a neighbourhood of a . U specifies those values everywhere, which is important, for example, when we are looking at the Lagrangian formulations of Newtonian mechanics.

reasons (see Albert [1]). However it is natural to figure out to what extent this picture commits us to configuration space substantivalism, and how far we can get without reifying configuration space. It turns out that it isn't that hard to do without it provided we are ordinary substantivalists about space-time. To get a quantity defined over ordinary space that will do the job, all one needs is a quantity that requires many arguments to determine its value, N spatial points to be precise, rather than a single point of configuration space. Formally this is a real function $U(p_1, \dots, p_N)$ taking N space-time points and returning a scalar quantity. For conceptual clarity, it's important not to confuse an N -place function with a 1-place function taking a single argument with N components to a real number (i.e. a scalar field on configuration space) — even though they are mathematically isomorphic, the latter formalism assumes a richer ontology of points of configuration space.

The second thing to note is that the potential field U is not a perspicuous way to represent reality since its values depend covertly, and in a non-perspicuous way, on the properties of three particular particles.⁴⁶ For example, in a system with three particles, a , b and c , the potential is completely described by all the facts of the form $U(p_1, p_2, p_3) = \alpha$, stating that when a is at p_1 , b at p_2 and c at p_3 , the potential is α . Clearly the value of $U(p_1, p_2, p_3)$ depends on the masses of a, b and c : as we vary the masses of a, b and c through modal space this value changes, whereas if we were to vary the masses of three distinct particles the value wouldn't change. Thus two qualitatively identical worlds could differ over the value of $U(p_1, p_2, p_3)$ if a, b and c existed and exerted forces in one world, and didn't exist (and thus didn't exert forces) in the other.⁴⁷

It is somewhat strange to think that the fundamental laws governing our example above are essentially specific to the motions and forces generated by the particular particles a, b and c , and not about how three particles in general would interact with one another. It should also strike us as unsatisfactory that in order to say what the potential is when three different particles with the same masses are located at $p_1 - p_3$ one must introduce a new fundamental primitive. These issues can easily be resolved if we adopt a more general relation as our fundamental primitive, that takes both particles and space-time points as arguments, and tells you what the potential is when those particles are located at those positions. In the three particle case, for example, we'd employ a six place function $U(a, p_1, b, p_2, c, p_3) = \alpha$ which tells us what the potential is when a is at p_1 , b at p_2 and c is at p_3 . Thus the three place function we were working with above is just the result of substituting three particular particles as arguments into our more fundamental function.

The third and final modification we need to make is to eliminate reference to numbers in our formulation of the potential. Following Field [11] we achieve this by means of congru-

⁴⁶A related point is made by Bradley Monton about configuration space realism in the context of quantum mechanics. See [22].

⁴⁷In a world where a, b and c do not exist, $U(p_1, p_2, p_3)$ is 0, as we have introduced it, since a, b and c do not exert any forces. One might try to come up with a different way of interpreting U so that its values somehow depend on whichever three particles happen to exist (assuming there are only three), but the hopes for that project seem dim. Consider the set of worlds, X , in which only the particles a, b and c exist in different possible arrangements. Now imagine two sets of worlds, $X[a/d]$ and $X[a/d']$, whose worlds are qualitatively identical to the worlds in X in which particles d and d' respectively have been substituted for a in each world in such a way that they play the some qualitative role as a (and assume b and c remain in their qualitative roles). We would presumably want the first argument of U to track the mass of d throughout $X[a/d]$ and d' throughout $X[a/d']$. By completely symmetrical reasoning, if we were to substitute b in each world in $X[a/d]$ with d' we'd expect the second argument of U to track the mass of d' throughout throughout the resulting space of worlds, $X[a/d][b/d']$. Similarly if we were to substitute b with d in each world in $X[a/d']$, we'd expect the second argument of U to track the mass of d in $X[a/d'][b/d]$. However this is impossible because $X[a/d'][b/d] = X[a/d][b/d']$ (since for every world in X there's a qualitatively identical world in X in which a and b have switched roles).

ence and betweenness relations. For example, assuming configuration space substantivalism we could adopt a four-place congruence relation $C_U(x, y, z, w)$ and three-place betweenness relation $B_U(x, y, z)$ on configuration space. The former tells us when the difference in potential between two points in configuration space is the same as between two other points, and the latter tells us when its value at a point is between the values at two other points respectively. As we noted above, we can reformulate this theory in a way that doesn't assume configuration space substantivalism, by introducing relations over ordinary points of space with more arguments. Thus congruence becomes a $4N$ -place relation $C_U(\bar{x}, \bar{y}, \bar{z}, \bar{w})$ and betweenness a $3N$ place relation, $B_U(\bar{x}, \bar{y}, \bar{z})$ where each barred variable is short for a sequence N independent variables standing for space-time points, rather than a single variable standing for a point in configuration space. The large number of arguments should not come as a surprise: if we are not assuming configuration space substantivalism, even our potential function U has to have a large number of argument places.⁴⁸ Finally, we noted that gravitational potential is most perspicuously represented by including the particles as arguments as well; $U(a_1, p_1, \dots, a_N, p_N)$. This function can be captured using betweenness and congruence relations in the usual way, although now we double the number of arguments.⁴⁹ General laws — sentences in which all the arguments of primitive predicates are occupied by bound variables — stated in this language will state qualitative facts.

5.2 Fields and the theory of multiple location

The framework we have just described is at least a gesture at what I take to be the most promising assortment of fundamental relations that can recover the fields of Newtonian physics without postulating fundamental relations between the physical and the abstract. What is particularly interesting, however, is that this formalism extends without any modification to the setting of interest to us, in which every object is located at every Galilean transform of its locations.

To get an intuitive feel for this, note that the fundamental relations we have outlined above determine (via a potential) the forces acting at each point in configuration space. For Newtonian systems the potential on configuration space has the following form:

$$U = \sum_{i < j} U_{ij}(|x_i - x_j|) = \sum_{i < j} Gm_i m_j / |x_i - x_j|$$

her i and j range over numbers less than the dimension of the configuration space (three times the number of particles), $x_{3k}, x_{3k+1}, x_{3k+2}$ describe the three components of the location of the k th particle, $m_{3k} = m_{3k+1} = m_{3k+2}$ is the mass of the k th particle, and G the universal gravitational constant.

The components of the slope of this field, recall, determines a force field on configuration space that states what the forces acting on each particle would be if the particles were configured according to that point. Recall also that space-time symmetries of the particles, such as translations and rotations, correspond to symmetries of this field in configuration space. In our example with two particles constrained to a single dimension, configuration space was two dimensional and the potential U at each point is, given our equation above, proportional to 1 over the distance of that point from the line $y = x$ (turning infinite at $y = x$, corresponding to the collision of the particles). The level sets of U (the set of points

⁴⁸The large number of arguments is a feature that was introduced when we rejected configuration space realism; it is not the use of congruence and betweenness relations themselves that are responsible for this.

⁴⁹If you're keeping track, congruence now has $8N$ arguments and betweenness $6N$, where N is the number of particles.

at which U has a constant value) thus consist of pairs of parallel lines that are reflections of each other about the diagonal line $y = x$. A translation of both particles in space would correspond to a diagonal shift in configuration space along a level set of U , and a reflection in space would take me from one diagonal line to the diagonal line on the other side of and of equal distance from the line $y = x$. (Note that there are no rotations for particles constrained to one dimension. When configuration space is higher dimensional the level sets of U tend to have a much more interesting structure). The trajectory two systems trace out in configuration will have the same shape if they start out with the same velocity and at points in configuration space belonging to the same level set.

A multi-particle system, consisting of singly located particles at various space-time points, corresponds to a single point of configuration space: the ‘systems location in configuration space’. According to our version of substantivalism, a physical system is multiply located at all of its Euclidean transforms. In configuration space this means the system is multiply located at every point in a given level set of U . We can now see how our original problem for space-time fields is avoided by moving to fields over configuration space. Given an ordinary field on space-time we cannot uniquely associate a field value to a multiply located particle: the value of the space-time field can be different at different locations. However note that a system that’s multiply located in accordance with SHIFT EQUIVALENCE is only located at points of configuration space where the value of U is the same. Thus by working with fields on configuration space — or at any rate, working with many place relations on space-time that determine such a field on configuration space — we no longer find ourselves associating multiple field values to the same particle.

6 Appendix B: the first-order theory of part and location

In this appendix we explore further some of the issues relating to a first-order axiomatisation of the pure theory of location and part. That theory relies on two definitions:

DEFINITION: Say that the sequence x_1, \dots, x_n is located (or 0-strongly located) at the sequence y_1, \dots, y_n , written $L^n x_1 \dots x_n y_1 \dots y_n$, iff the fusion of any elements of the former sequence is located at the fusion of the corresponding elements of the latter sequence.

DEFINITION: A sequence x_1, \dots, x_n of material objects is a *partition* if its elements are pairwise disjoint but their fusion is the fusion of all material objects. A subpartition of length $m > n$ of $x_1 \dots x_n$ is a partition $y_1 \dots y_m$, such that every x_i is a fusion of a subsequence of $y_1 \dots y_m$.

Note that for a fixed n these definitions can be formulated in first order logic. In the former case, for example, it would involve a long conjunction with a conjunct for all the possible subsequences of $1, \dots, n$.⁵⁰

With these definitions at hand we can state a finitary version of ARBITRARY PARTITIONING in first order logic (the thesis is schematic, for each choice of n):

⁵⁰The first definition is given by the formula:

$$L^n x_1 \dots x_n y_1 \dots y_n := \bigwedge_{\{i_1 \dots i_k\} \subseteq \{1 \dots n\}} Lx_{i_1} + \dots + x_{i_k} y_{i_1} + \dots + y_{i_k}$$

PARTITIONING: If x is located at y , then there is some y' disjoint from y such any n -length subpartition of x, x^c is located at some n -length subpartition of y, y' .

Here x^c denotes the complement of x in the material objects: the fusion of material objects disjoint from x . Roughly what the axiom says is that if x is located at y , and we cut up x and x 's complement into lots of little bits, we can find a way of cutting up y and y 's complement (relative to some region) in such a way that not only are each of the former bits located at the latter bits, but fusions of the former bits are located at corresponding fusions of the latter bits.

If we add to the above axioms the principle that every object has a unique location we can recover the classical theory of location and part. In particular, we can show that the location function preserves fusions and relative complements.⁵¹

However unfortunately the above axioms (along with the other first-order axioms of section 3.3) are not enough to pin down the structure of the location relation. There are relations that satisfy the axioms but are not unions of location functions. Part of the issue here is that this isn't even possible in the first-order versions of the classical theory of location and part: those axioms require that location functions preserve arbitrary fusions, which is not something you can state in a first order language.⁵² There are, however, a couple of further questions that bear investigation:

QUESTION 1: Say that a function from a domain of material objects to a domain of regions is a weak location function if it preserves atomhood, relative complements and *finitary* fusions. Is it true that every relation satisfying our axioms is a union of weak location functions?

QUESTION 2: Is our axiomatisation complete? Are there any principles that are true in all the intended models that are not provable from our system? (Here we may consider the class of all the models in which L is a union of location functions, or the class of models where L is a union of weak location functions, as our intended models.)

Although both of these questions are open, there is a sequence of principles that at least *prima facie* seem to strengthen PARTITIONING; if they are indeed stronger then the answer to the second question is negative. These principles require another definition:

DEFINITION: A partition x_1, \dots, x_n is $k + 1$ -strongly located at y_1, \dots, y_n if and only if for every m -length subpartition of x_1, \dots, x_n is k -strongly located at some m -length subpartition of y_1, \dots, y_n .

⁵¹To see this note that $x \sqcap y, x \setminus y, y \setminus x$ (assuming they all exist (if they don't all exist, the following argument will simplify) is a partition of $x + y$, which by PARTITIONING means that there is a partition of $x + y$ s location, r, s, t , such that the sequence $x \sqcap y, x \setminus y, y \setminus x$ is located at r, s, t . This means in particular that $x = x \sqcap y + x \setminus y$ is located at $r + s$, and $y = x \sqcap y + y \setminus x$ is located at $r + t$, that $x \setminus y$ is located at $s = (r + s) \setminus (r + t)$, which is $l(x) \setminus l(y)$, so that our location function preserves relative complements. Similarly, since x is located at $r + s$, y at $r + t$, and $x + y$ at $r + s + t$ we see that the location of $x + y$ is just the fusion of the location of x with the location of y , demonstrating that our location function preserves fusions. Thus we can see that the above theory is a generalisation of the classical theory of location and part.

⁵²Let our domain of material objects be \mathbb{N} . Let X be the quotient Boolean algebra $P(\mathbb{N})/\approx$, where $A \approx B$ iff $|A \setminus B \cup B \setminus A|$ is finite. For A in $P(\mathbb{N})$, let $[A]$ be the member of X containing A . By Stone's representation theorem X is isomorphic to some algebra of subsets of Y (disjoint from \mathbb{N}), by some isomorphism f . Let the set of points of space be $\mathbb{N} \cup Y$. The location function $l(A) = A \cup f([A])$ satisfies all of the axioms above but does not preserve arbitrary fusions since $l(\{n\}) = \{n\}$ for each n but $l(\mathbb{N}) = \mathbb{N} \cup f([\mathbb{N}])$. Thanks to [REF] for suggesting this construction.

Here, of course, our definition is encoded in a different first-order formula for choices of k , m and n , where $m > n$. As m and n increase the number of conjunctions in the first order definitions of ‘located’, ‘partition’ and ‘subpartition’ increase, and as k increases the number of alterations of the quantifiers increases as well. The case $k = 0$ is given by our original notion of location for sequences. We can thus strengthen PARTITIONING to:

k -PARTITIONING: If x is located at y , then for some y' disjoint from y , x, x^c is k -strongly located at y, y' .

Once again, if these principles are not provable from our initial system, it is natural to ask questions analogous to the two questions above about a strengthened system containing these principles.

References

- [1] David Albert. Elementary quantum metaphysics. In J. T. Cushing, Arthur Fine, and Sheldon Goldstein, editors, *Bohmian Mechanics and Quantum theory: An Appraisal*, pages 277–284. Kluwer, 1996.
- [2] Frank Arntzenius. *Space, Time, & Stuff*. Oxford Univ. Press, 2012.
- [3] Frank Arntzenius and Cian Dorr. *Calculus as Geometry*, chapter 8, pages 213–268. Oxford University Press, 2012.
- [4] Julian B. Barbour and Bruno Bertotti. Gravity and inertia in a machian framework. *Nuovo Cimento*, 38:1–27, 1977.
- [5] Julian B. Barbour and Bruno Bertotti. Mach’s principle and the structure of dynamical theories. *Proceedings of the Royal Society, London*, pages 295–306, 1982.
- [6] Gordon Belot. *Geometric Possibility*. Oxford University Press, 2011.
- [7] Roberto Casati and Achille C. Varzi. *Parts and Places. The Structures of Spatial Representation*. The MIT Press, 1999.
- [8] Shamik Dasgupta. Individuals: An essay in revisionary metaphysics. *Philosophical Studies*, 145(1):35–67, 2009.
- [9] Cian Dorr and Jeremy Goodman. Diamonds are forever. *Nous*, Forthcoming.
- [10] M. Eddon. Fundamental properties of fundamental properties. In Karen Bennett Dean Zimmerman, editor, *Oxford Studies in Metaphysics, Volume 8*, pages 78–104. 2013.
- [11] Hartry Field. *Science Without Numbers*. Princeton University Press, 1980.
- [12] Hartry Field. Can we dispense with space-time? *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1984:33–90, 1984.
- [13] Kit Fine. The non-identity of a material thing and its matter. *Mind*, 112(446):195–234, 2003.
- [14] Graeme Forbes. Time, events, and modality. In Robin Le Poidevin and Murray MacBeath, editors, *The Philosophy of Time*, pages 80–95. Oxford University Press, 1993.

- [15] David Hilbert. *The foundations of geometry*. Open court publishing Company, 1899.
- [16] Shieva Kleinschmidt. Multilocation and mereology. *Philosophical Perspectives*, 25(1):253–276, 2011.
- [17] Shieva Kleinschmidt. Shaping up location: Against the humean argument for the extrinsicity of shape. *Philosophical Studies*, 172(8):1973–1983, 2015.
- [18] Gottfried Wilhelm Leibniz and Samuel Clarke. The leibniz-clarke correspondence. In Elizabeth Schmidt Radcliffe, Richard McCarty, Fritz Allhoff, and Anand Vaidya, editors, *Late Modern Philosophy: Essential Readings with Commentary*. Blackwell Pub. Ltd., 2007.
- [19] Tim Maudlin. Buckets of water and waves of space: Why spacetime is probably a substance. *Philosophy of Science*, 60(2):183–203, 1993.
- [20] Tim Maudlin. Time, topology and physical geometry. *Aristotelian Society Supplementary Volume*, 84(1):63–78, 2010.
- [21] Tim Maudlin. *Philosophy of Physics*. Princeton Up, 2012.
- [22] Bradley Monton. Quantum mechanics and 3ndimensional space. *Philosophy of Science*, 73(5):778–789, 2006.
- [23] Brent Mundy. The metaphysics of quantity. *Philosophical Studies*, 51(1):29–54, 1987.
- [24] Oliver Pooley. Substantivalist and relationalist approaches to spacetime. In Robert Batterman, editor, *The Oxford Handbook of Philosophy of Physics*. Oxford University Press, 2013.
- [25] Jeffrey Sanford Russell. On where things could be. *Philosophy of Science*, 81(1):60–80, 2014.
- [26] Raul Saucedo. Parthood and location. In Karen Bennett and Dean Zimmerman, editors, *Oxford Studies in Metaphysics Vol. 6*. Oxford University Press, 2011.
- [27] Theodore Sider. *Four Dimensionalism: An Ontology of Persistence and Time*. Oxford University Press, 2001.
- [28] Bradford Skow. Are shapes intrinsic? *Philosophical Studies*, 133(1):111–130, 2007.
- [29] Alfred Tarski. What is elementary geometry. *The axiomatic Method, with special reference to Geometry and Physics*, pages 16–29, 1959.
- [30] Alfred Tarski. *Logic, Semantics, Metamathematics Papers From 1923 to 1938*. Hackett, 1983.
- [31] Gabriel Uzquiano. Mereological harmony. In Karen Bennett and Dean Zimmerman, editors, *Oxford Studies in Metaphysics Vol. 6*. Oxford University Press, 2011.