

WORLDLY INDETERMINACY: A ROUGH GUIDE

Gideon Rosen, Princeton University and
Nicholas J.J. Smith, Victoria University of Wellington

ABSTRACT: This paper defends the idea that there might be vagueness or indeterminacy in the world itself — as opposed to merely in our representations of the world — against the charges of incoherence and unintelligibility. First we consider the idea that the world might contain vague *properties and relations*; we show that this idea is already implied by certain well-understood views concerning the semantics of vague predicates (most notably the fuzzy view). Next we consider the idea that the world might contain vague *objects*; we argue that an object is indeterminate in a certain respect (colour, size, etc.) just in case it is a borderline case of a maximally specific colour (size, etc.) property. Finally we consider the idea that the *world as a whole* might be indeterminate; we argue that the world is indeterminate just in case it lacks a determinate division into determinate objects.

There is a widespread suspicion that indeterminacy and vagueness can only be features of claims or representations, and hence that it makes no more sense to ask whether the world is indeterminate than it does to ask whether the world rhymes or whether it's written in English. On this view, it is not a genuine possibility that the world itself — as opposed to our representations of it — might be unsettled or inchoate in some way: the world must be a fully determinate array of facts or things.¹

Our aim in this paper is to make clear sense of the idea that the world itself might be vague or indeterminate in some respect. We shall not be arguing that the world *is* indeterminate. We shall not even be arguing that worldly indeterminacy is a genuine metaphysical possibility. Our aim is the more basic one of showing that

¹ This was David Lewis's view:

The only intelligible account of vagueness locates it in our thought and language. The reason it's vague where the outback begins is not that there's this thing, the outback, with imprecise borders; rather, there are many things, with different borders, and nobody's been fool enough to try to enforce a choice of one of them as the official referent of the word 'outback'. [Lewis 1986: 212]

We disagree. And yet it goes without saying that our thinking on this topic (as on most others) owes an incalculable debt to David's work.

contrary to the widespread suspicion just mentioned, the idea cannot be ruled out on grounds of incoherence or unintelligibility.

VAGUE PROPERTIES AND RELATIONS

Suppose that a theory of vague language (specifically, of vague predicates) tells us that the sentence 'Bob is tall' is neither simply true nor simply false. (Perhaps it lacks a truth value; perhaps it has an intermediate truth value.) Does that mean that the world contains a vague property, *tallness*? That depends, in part, on the semantics for vague sentences.

Consider the standard supervaluationist view.² Suppose 'Bob is tall' comes out true on some admissible precisifications of 'tall' and false on others. Then 'Bob is tall' is neither true nor false. But this does *not* mean that the property of tallness is vague. The point is rather that there is no unique property of tallness. There are many properties, each corresponding to a sharpening of 'tall', and all precise, in the sense that each object either definitely possesses or definitely fails to possess each property. When we speak vaguely we fail to single out a unique such property. So vagueness is solely a matter of the relation between language and the world. The world itself is precise.

On the standard fuzzy view, on the other hand, predicate vagueness is primarily a metaphysical matter. When I say 'Bob is tall' I pick out a unique property, *tallness*: there is no semantic indecision here at all (i.e., no vagueness in the relationship between language and the world). Rather, the property I pick out (which corresponds to the fuzzy set of tall things) is such that things can possess it to any degree between 0 and 1. Thus the reason my sentence is neither simply true nor simply false is that the predicate picks out a unique property of which Bob is a borderline instance. It is a matter of how things are in the world that Bob neither definitely possesses nor definitely fails to possess the property *tallness*.

The claim that the fuzzy theory involves a commitment to vague properties and relations while the supervaluationist theory does not involves 'taking the model theory literally'. But how else should we take it? Fuzzy and classical model theory agree that each discourse is associated with a unique privileged interpretation: the *intended* or *correct* interpretation. (The latter term is due to [Islam 1996].) This

interpretation has a domain, and it assigns a function from the set of n -tuples of members of the domain to a set of truth values as the interpretation of each n -place predicate. Where the accounts differ is over the set of truth values. The classical theory takes it to be the set $\{0,1\}$, while the fuzzy theory takes it to be the interval $[0,1]$. The supervaluationist departs from the classical view in a different way: her interpretations are classical, but instead of associating one privileged interpretation with each discourse, she associates many admissible interpretations, with supervaluationist truth (or supertruth) being classical truth on all admissible interpretations. Now the fuzzy model theory gives us a picture in which each n -place predicate is associated with a unique n -place relation, while the supervaluationist model theory gives us a picture on which each n -place predicate is associated with many n -place relations (one in each admissible interpretation), each of which is inherently precise. Now if we think the fuzzy view is the correct semantic theory for vague language, we think that when we use a vague predicate we refer to a particular relation which is inherently vague; and if we think the supervaluationist view is correct, we think that when we use a vague predicate, we do not refer to a particular relation, but all the candidate relations we simultaneously refer to are inherently precise. So if the fuzzy view is correct, there exist vague properties and relations alongside the precise ones.³ And this is already one sort of metaphysical vagueness.

Is the fuzzy view correct? Our task here is to make sense of metaphysical vagueness, not to argue that it exists. Thus we need not make the case that the fuzzy view is correct. It is enough that it makes sense — and this is certainly the case.⁴

² That is, a supervaluationist view on which the admissible precisifications must be classical.

³ Indeed the precise ones are those special cases of the vague ones that map all n -tuples to 0 or 1.

⁴ It is worth stressing that hybrid views are possible, and indeed plausible. The essential feature of the fuzzy view is its incorporation of fuzzy semantic values. Fuzzy theorists standardly assume a single correct assignment of these fuzzy values to the non-logical words. But there is no reason in principle why he should not allow for semantic indecision. It might be that there are many equally admissible fuzzy interpretations of the language, or even that the many candidates vary in their degree of admissibility. The fuzzy semantics for such a language might then take a leaf from the supervaluationist's book. (For discussion of a view along these lines, see [Smith forthcoming: §13].) In what follows we shall often write as if each vague predicate is associated with a single fuzzy value. But that is a simplifying assumption. We take no view on whether it is true, nor even on whether it could possibly be true.

VAGUE OBJECTS

The fuzzy theory presupposes that the world contains vague properties and relations. But when we hear the suggestion that the world itself might be vague, we normally take it to be the claim that the world contains vague objects: not just fuzzy properties but fuzzy *things*. What might it mean for an object to be vague or indeterminate in some respect? One natural idea is that an object is vague when it is a vague instance — a borderline case — of some vague property. This thought is encouraged by an array of compelling examples. For example, suppose we have an object, say Mt. Everest, with the following feature: there is a spacetime point p such that it is neither determinately the case nor determinately not the case that p falls within its boundaries. If this is a matter of worldly vagueness of the location relation (as opposed to mere semantic vagueness of the location predicate or the name 'Mt. Everest'), then it is natural to say that Mt. Everest is in one sense a vague object. Or again, suppose we had an object, say a cloud, with the following feature: it definitely weighs more than one ton and less than two tons, but it is neither determinately the case nor determinately not the case that it weighs exactly 1.36 tons. If this is a matter of worldly vagueness — if there is no semantic indecision in the picture — then it is natural to say that our cloud is in one sense a vague object.

Such examples suggest the simple thought that an object is vague when it is a borderline case of a vague property. But of course this can't be true in general. Let B be a small turquoise bead, perfectly uniform in colour, and suppose, as is plausible, that the predicate '... is blue' is neither determinately true nor determinately false of it. (Turquoise is too green to be clearly blue, and too blue to be clearly green.) Suppose further that the predicate '... is blue' corresponds to a single vague property, *blue*, and hence that B is a borderline case of this property. Intuitively, this has no tendency to show that B is indeterminate in colour, nor indeed in any respect. To the contrary, the case invites us to imagine that B is a perfectly determinate shade of colour through and through.

So not every property makes a vague object out of its borderline cases. Some apparently do — e.g., the location and mass properties mentioned above; and some apparently don't — e.g., *blue*. And that sets our problem. If we had a way of distinguishing the properties that make vague objects out of their vague instances — call them the *sharp* properties — then we could refine the natural thought. We could say that an object is vague just in case it's a borderline case of a sharp property. The challenge would then be to say which properties count as sharp.

Our bead B is not genuinely indeterminate in colour despite being a borderline instance of *blue*. But suppose we told you that we had another bead B^* that really was indeterminate in colour. What might you take us to have in mind? Something like this, we imagine. B^* might be a clear case of some colour properties. It might be clearly blue, as opposed to red or green. But when it came to saying which shade of blue it was, the question would have no clear answer. The colours are standardly represented by a three-dimensional solid whose points correspond to particular shades — what an older tradition would have called the *determinates* of the various *determinable* colour properties. *Blue* corresponds to some sort of fuzzy region in this solid; and you don't get to be indeterminate in respect of colour simply by sitting on the penumbra of a fuzzy region. When we say that B^* is indeterminate in respect of colour, we can only mean that it is a borderline case of one or more *particular shades* of colour: definitely blue, perhaps, but neither definitely $B-17$ nor definitely not $B-17$, where ' $B-17$ ' is the decorator's name for some *point-sized* region in colour space. It seems that insofar as we have an intuitive conception of what indeterminacy in colour might amount to, it must be something along these lines.

The idea generalizes. If we told you that we had an object that was genuinely indeterminate in size or mass, you should take us to mean that there was no particular or determinate mass property or size property that the thing determinately possessed. Our aim in what follows is to present an explicit formulation of this idea.

POINT PROPERTIES

We shall assume that properties come in *families* or *categories*. The colours, the masses, the shapes, the volumes, the chemical structures, the arrangements of component atoms — all of these plainly constitute, in some intuitive sense, relatively maximal families or categories of properties. A relatively maximal family contains, for each of its members, every other property of the same *kind*. It will contain the *determinables* as well as the *determinates* that fall under them. *Triangularity* is a shape property, so it falls in the category of shape. But a triangular figure may have any number of shapes, and each of these properties is also a shape property in the category of shape. We won't attempt an analysis of this notion of a category. But we stress that the partial inventory of folk categories listed above is entirely provisional. We take it to be a broadly empirical matter — and in any case a matter for substantive metaphysics — what sorts of properties there are and how the lines in property space are to be drawn. What our analysis presupposes is that no matter how this turns out,

there will be a distinction between relatively complete or maximal families of properties and the rest. If this notion breaks down, the account of worldly indeterminacy will have to proceed along very different lines.

What we need is a general account (one that works for all categories) of the difference between properties such as *blue* which do not make indeterminate objects out of their vague instances, and properties such as *B-17* which do. The crucial difference appears to be one of specificity: settling that an object is blue still leaves open a lot about its colour; not so, settling that it is *B-17*. Let's call properties of the sort we're interested in **point properties**. Our suggestion is that for an object to be intrinsically indeterminate in respect of some category is for it to be an intermediate instance (a borderline case) of a point property in that category, and that an object is indeterminate *sans phrase* when it is indeterminate in respect of some category. The remaining question is how the general notion of a point property is to be captured.

One obvious constraint is that point properties must be intrinsic. A bead B^* may be indeterminate in colour if it is an indeterminate instance of *B-17*. But the moon should not turn out to be an intrinsically indeterminate entity just because it is a borderline case of *being such that B^* is B-17*.⁵

We shall assume that each category F is associated with a relation of exact similarity, \approx_F : x is exactly the same colour (shape, mass, length) as y . For the sake of having a relatively concrete framework to work in, we shall also suppose that properties are (modelled by) fuzzy subsets of the domain D of all possible objects, and that n -place relations are (modelled by) fuzzy sets of n -tuples on D . We shall use the notation $P(x_1 \dots x_n) = r$ to mean that the n -tuple $(x_1 \dots x_n)$ is a member of the relation P to degree r , $0 \leq r \leq 1$. Thus, for example, $Red(\text{Tom}) = 0.7$ means that Tom (the tomato) possesses the property *red* to degree 0.7.⁶

⁵ This requirement threatens the location example, because location is presumably not an intrinsic property. However, the essential idea behind the example can easily be repackaged in terms of *extent* or *size* rather than *location*. Thus, even restricting ourselves to intrinsic properties, Mt Everest might still turn out to be a genuinely vague object, if it is a borderline case of some specific size property.

⁶ We stress that this appeal to the fuzzy framework is just for the sake of concreteness. Nearly everything we say will go over to the simple three-valued case or to more complicated many-valued approaches such as the approach defended in [Smith forthcoming].

We shall not assume that for every x and y , either $\approx_F(x, y) = 1$ or $\approx_F(x, y) = 0$. Nevertheless, \approx_F is the relation of *exact* similarity in respect of F . The idea is this. If $\approx_F(x, y) = 1$, then x and y are definitely exactly similar in respect of F . If $\approx_F(x, y) = 0$ then they are definitely not exactly similar in respect of F — but they might still be extremely similar. A bicycle A painted bright yellow, and another bicycle B painted with the same paint after the addition of a drop of orange tinter will be *extremely* similar in respect of colour; but they are definitely not exactly similar. They will be in principle distinguishable, after all, if only by indirect means. So in this case, $\approx_{\text{colour}}(A, B) = 0$. On the other hand, if $\approx_{\text{colour}}(x, y)$ is strictly between 0 and 1, then it is neither definitely the case that x and y are exactly similar, nor definitely the case that they are not exactly similar. Suppose, for example, that x is a definite case of $B-17$, but that y is a borderline case. Then since it's not fully determinate which colour y has, it's not fully determinate whether x and y are exactly similar in respect of colour.

The basic thought about point properties is that they are maximally specific: hence if two objects are degree-1 instances of a point property P , they must be definitely exactly similar in the relevant respect. So our first condition that a point property P in category F must satisfy is as follows.

(a) (Necessarily) If $P(x) = P(y) = 1$, then $\approx_F(x, y) = 1$.

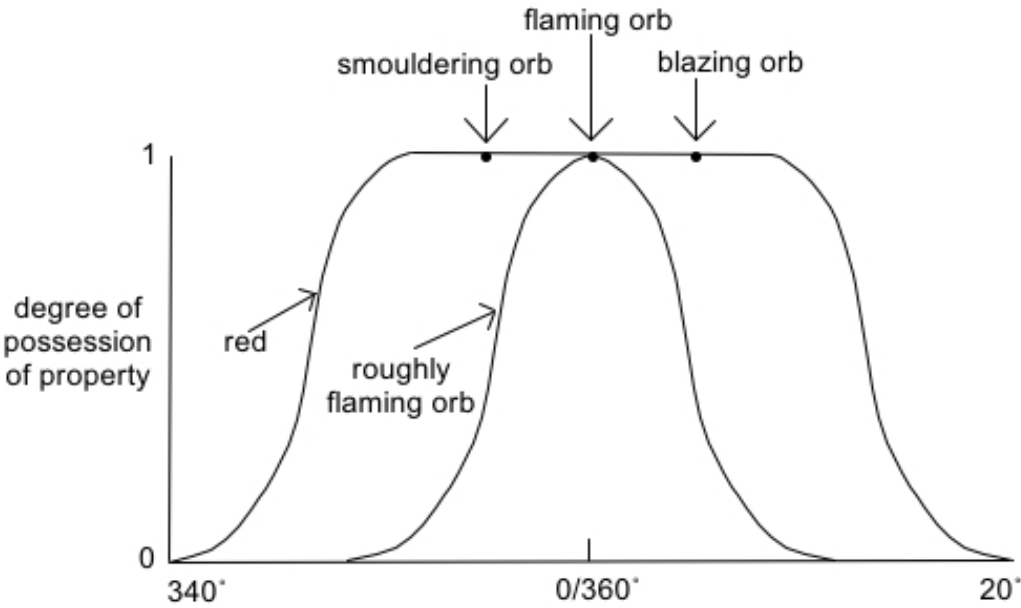
Consider the category of colour properties. Let's suppose that *flaming orb* is a maximally determinate shade of red, corresponding to a particular point on the colour solid. If two chairs are both painted flaming orb, they will be exactly similar in colour. So *flaming orb* meets our first condition. On the other hand, *red* clearly flouts it. The ripe tomato and the pillar box may both be degree-1 instances of *red* while nevertheless differing somewhat in shade.

Any property that cannot have degree-1 instances will satisfy our first condition, but intuitively such properties will not in general be point properties. Consider the property of *being red and orange*. Plausibly, it has no degree-1 instances. But it should not count as a point property. For recall our basic idea: being an indeterminate instance of a point property makes an object indeterminate in respect of F . This autumn leaf may be a borderline case of red, and a borderline case of orange, and hence (given some natural assumptions) a borderline case of *red and orange*. But it is not indeterminate in respect of colour: it may have a perfectly definite

colour somewhere between focal red and focal orange. So we add a second condition:

(b) (Possibly) there is an x such that $P(x) = 1$.

It is natural to suppose that conditions (a) and (b) together capture the notion of a point property; but they don't, and this is where things get interesting. Consider the property we might call *roughly flaming orb*. The degree-1 instances of this property are exactly the degree-1 instances of *flaming orb*. However, whereas an object which is a slightly different shade from *flaming orb* (say *smouldering orb*) is a degree-0 instance of *flaming orb*, no matter how close the shades may be, *roughly flaming orb* works differently. Possession of this property trails off gradually the further one gets from *flaming orb*, so that a degree-1 instance of *smouldering orb* is (say) a degree 0.5 instance of *roughly flaming orb*. (See figure.⁷)



⁷ We are on the 'equator' of HLS space (i.e., with lightness at 50% and saturation at 100%), considering variation in hue on the horizontal axis of the diagram. The essential features of the diagram are (i) that *red* has a relatively wide region of degree-1 instances, while the other properties have only point-sized regions of degree-1 instances; (ii) that the degree-1 points of *flaming orb* and *roughly flaming orb* coincide, while the degree-1 points of *smouldering orb* and *blazing orb* are distinct but nearby; (iii) that the only non-zero points of *flaming orb*, *smouldering orb* and *blazing orb* are degree-1 points, while *red* and *roughly flaming orb* have penumbral regions, and (iv) that the penumbral region of *roughly flaming orb* includes the degree-1 point of *smouldering orb*.

Now *roughly flaming orb* meets our conditions. It has degree-1 instances, and its degree-1 instances are all exactly similar in colour. And yet it should not count as a point property given the use we want to make of that notion. For objects that are perfectly determinate in colour — the degree-1 instances of *smouldering orb*, for example — are intermediate instances of *roughly flaming orb*. So the question arises: how are properties of this sort to be excluded?

It will not do to say that *roughly flaming orb* fails to be a point property because it has vague boundaries. That would spoil our main hypothesis, viz., that objects can be vague in respect of (say) colour, and that their vagueness consists in being borderline instances of a point property in the relevant family. If this is right then *flaming orb* — our paradigm case of a point property — will turn out to have borderline cases as well.

Of course there is this difference. The borderline cases of *roughly flaming orb* may be perfectly determinate in colour: not so the borderline cases of *flaming orb*. But however true this may be, it won't do as a definition for our purposes. Our aim is to break out of this circle by defining the notion of a point property in F independently, so that determinacy in respect of F may be defined in terms of it.

So far as we can see, the distinction between genuine point properties like *flaming orb* and pseudo-point properties like *roughly flaming orb* cannot be drawn in terms of the materials we have so far developed. We need a new idea. So let's consider why we might think that *flaming orb* intuitively counts as a point property in a sense in which *roughly flaming orb* does not. We begin with the observation that there is a palpable sense in which *roughly flaming orb* is a derivative property. If we ask what it is for a thing to be an instance of *roughly flaming orb*, the answer will presumably run as follows. To be a degree-1 instance of *roughly flaming orb* is to be a degree-1 instance of *flaming orb*; to be a degree 0.5 instance of *roughly flaming orb* is to be a degree-1 instance of *smouldering orb* (or *blazing orb*) etc. The full story is told in the graph. On the other hand, if we ask for a philosophical account of what it is for a thing to be an instance of *flaming orb*, it's not clear what we could say. We might say, 'Well, it is in part to be red'. Or we might say, 'It is to reflect light of such and such a wavelength'. But in either case it seems clear that the account of what it is to possess *flaming orb* — the real definition of the property, as it were — will not advert to *roughly flaming orb*. Indeed it is plausible that it will not advert to any other similarly determinate colour property.

This is an asymmetry. But what exactly does it come to? The notion of real definition to which we appeal is relatively familiar, if not altogether uncontroversial. We do not know how to explain it in more basic terms, and we do not pretend to understand it fully. And yet we are inclined to regard it as intelligible, even in the absence of a worked out theory. Consider the following familiar thought. Goodman's *grue* and *bleen* are standardly defined in terms of *green* and *blue* and some third notion (say, *observed*) [Goodman 1983: 74]. But as Goodman shows, we can just as well define *green* and *blue* in terms of *grue* and *bleen* and the same third notion. And yet there is obviously a sense in which such a definition gets things the wrong way round. We might try to understand this in semantic terms, or in epistemic terms. But we might also try to understand it in metaphysical terms. It may be true that as a matter of necessity, x is green iff x is grue and observed or bleen and unobserved. But if we ask what it is to be green — if we ask what a thing's being green *consists in* — then it is just plain obvious that this extensionally correct verbal definition will not do. *Green* may be a disjunctive property, but it is not disjunctive in this way. The account of what it is for a thing to be green might be given in terms of reflectance profiles or by listing the particular shades of green and taking their disjunction. But whether a thing is observed has *nothing whatsoever* to do with whether it's green or why it's green.⁸ In our idiom this amounts to the claim that the property of being observed does not figure (in this way) in the account of what it is for a thing to be green (i.e., the real definition of the property *green*).

Much of what goes on in philosophy is the search for real definitions. When we seek a philosophical account of the colours, or of personal identity, or of the physical, we're not interested in any old counterexample-proof verbal definition. We want to know what it is for a thing to be red, or to be a person, or to be a physical entity. This is typically not a question about words or concepts. It is a question about the properties and kinds themselves. Sometimes we can begin to answer it a priori, but sometimes we can't. When the philosopher says [Johnston 1987] that to be a human being is to be an organism that is normally constituted by a human animal but which can survive reduction to the condition of a mere functioning brain, his account is obviously partly empirical. It is an a posteriori matter that the brain plays

⁸ Even an account of colour according to which *green* is the disposition to cause experiences of a certain kind in normal perceivers under standard conditions does not say that whether a particular object is green depends on whether that particular object is observed.

a role in mentation; it is an a posteriori matter that we have brains at all. So he's not spelling out the linguistic meaning of the word or the content of our concept of a human being. And yet his account cannot be dismissed on this ground. It might still be an adequate account of *what it is to be a human being*.

The notion of real definition nonetheless cries out for explication. We would like to know more about its formal properties, and about our epistemic access to its extension. But one way to clarify the notion while we await these desirable developments is to put it to use in the investigation of other topics. At any rate, that is what we propose to do here.

How does the notion help with our problem? The challenge is to distinguish genuine point properties from impostors like *roughly flaming orb*. And here our thought is as follows. The impostors will all be derivative properties: their definitions will advert to bona fide point properties, but not vice versa. The account of roughly flaming orb will have a clause specifying what it takes to be a degree-1 instance of the property, and this clause will take something like the following form: To be a degree-1 instance of the impostor is to be a degree-1 instance of another property. We therefore propose to solve the problem by excluding properties of this sort.

Let's say that a property P in family F is **point-like** iff it satisfies conditions (a) and (b) above: It has degree-1 instances, and its degree-1 instances are exactly similar in respect of F . *Flaming orb* and *roughly flaming orb* are both point-like, but only *flaming orb* satisfies the following further condition:

- (c) There is no point-like property $P^* \neq P$ in F such that for x to be a degree-1 instance of P just is for it to be a degree-1 instance of P^* .

The idea is that the point properties in F are the point-like properties whose real definitions do not depend (in a certain way) on other point-like properties in F .

PROLIFERATING POINT PROPERTIES

The account of point properties in the previous section was designed to include maximally specific shades such as *flaming orb* and to exclude both unspecific properties such as *red* and hybrid monsters such as *roughly flaming orb*. It turns out, however, that certain properties besides the ordinary precise shades meet our conditions on point properties.

Let's suppose that Tom is indeterminate in respect of colour because he is a degree-0 instance of all of the ordinary determinate shades except for *flaming orb* and *smouldering orb*, of each of which he is a degree-0.5 instance. Tom is definitely not (say) vermilion; but it's genuinely unsettled whether he is an instance of *flaming orb* or *smouldering orb*. Now consider the property Q which satisfies the following condition:

$$\begin{aligned}
 Q(x) &= 1 \text{ if } \textit{flaming orb}(x) = \textit{smouldering orb}(x) = 0.5, \text{ and for all other} \\
 &\text{ordinary shades } P, P(x) = 0 \\
 &= 0 \text{ otherwise.}
 \end{aligned}$$

Q may be thought of as Tom's **colour profile**: a complete account of how Tom stands with respect to the colour solid. Now there is a sense in which Q is itself a colour. It is a way of being coloured, indeed a maximally specific way. Intuitively, the colour of a thing (how it stands with respect to colour) is completely settled once we are told that it is Q in a way in which it is not settled when we are told that it is (say) red.

It turns out that Q is (by our definition) a point property in the category of colour. Its degree-1 instances stand in precisely the same relation to all of the ordinary determinate shades, and this presumably guarantees that they resemble one another exactly in respect of colour. (They are all perfect colour duplicates of Tom.) Moreover, the account of what it takes to be a degree-1 instance of Q does not invoke degree-1 possession of any other point-like properties. It thus appears that alongside the maximally specific determinate shades (lengths, masses), certain profile properties or **superpositional** properties — the maximally specific *indeterminate* shades (lengths, etc.) — will count as point properties in our sense.

Of course we could modify our definition so as to exclude Q and its ilk. Instead of (c) we could impose a more stringent condition:

(c*) There is no point-like property $P^* \neq P$ in F such that for x to be a degree-1 instance of P just is (in whole or in part) for it to be a degree- r instance of P^* .

On this sort of account, the real point properties are the point-like properties whose definitions do not make reference to any other point-like properties in the relevant

way.⁹ This account captures the sense in which the superpositional properties are second rate. They depend for their definitions on a more fundamental stratum of point properties.

The original account may be thought to retain certain virtues. As we have stressed, there is a sense in which Tom is in a perfectly determinate *colour state*, just as a particle whose quantum state is a certain superposition of *x spin up* and *x spin down* is nonetheless in a perfectly determinate quantum state. Our thinking about worldly vagueness exhibits a certain oscillation. On the one hand, when we think of an object like Tom as somehow poised between two perfectly determinate ordinary shades, we are inclined to say that Tom's colour state is somehow unsettled. On the other hand, once we entertain the possibility of an object that stands to the ordinary colour solid in this way, we are tempted to say that we have discovered a new class of colours — a new class of perfectly determinate ways of being coloured — much as quantum mechanics may be taken to have disclosed an utterly new class of ways of being located in space or possessing momentum. But then on the third hand it seems that these newly disclosed 'shades' are somehow derivative, and that the *real* determinates are the points in the ordinary colour solid, the ordinary determinate lengths, etc. Our two definitions capture the poles of this oscillation.

However, interesting as this may be, for present purposes we find a decisive reason to stop oscillating, and adopt condition (c*) rather than condition (c). Note for a start that even if we adopt condition (c) and hence allow the superpositional property *Q* defined above to count as a point property, Tom still counts as indeterminate in respect of colour. It is enough for this that Tom be a degree-0.5 instance of *flaming orb*. It makes no difference that Tom is also a degree-1 instance of some other point property in the category of colour. We would have a problem only if some object that should *not* count as indeterminate turned out to be an intermediate instance of some point property. We have not yet seen such a case. But consider the property *Q** which satisfies the following condition:

⁹ It would probably be a mistake to require that the definition of a point property never make reference to any other point property. In certain quantity domains it may be that the individual point properties — *being exactly 3 mm long, having a temperature of 40 degrees Kelvin* — are defined in terms of special point properties: the natural zero of the scale, or the natural unit, for example. These ways in which the definition of one point property can involve another do not fit the pattern sketched above, and we do not intend (c*) to exclude them.

$Q^*(x) = 1$ if *flaming orb* (x) = *smouldering orb* (x) = 0.5, and for all other ordinary shades P , $P(x) = 0$
 $= 0.5$ otherwise.

(The difference between the definitions of Q^* and Q is that the last '0.5' in the definition of Q^* was a '0' in the definition of Q .) Now a degree-1 instance of *flaming orb* — an object which intuitively *is* determinate in respect of colour — is a degree 0.5 instance of Q^* . Yet Q^* satisfies conditions (a), (b) and (c). However, Q^* does *not* fit condition (c*). The account of what it takes to be a degree-1 instance of Q^* invokes degree- r possession of certain other point-like colour properties.

Are there analogous counterexamples to an account which involves (a), (b) and (c*) rather than (c)? Consider the property Q^{**} which satisfies the following condition:

$Q^{**}(x) = 0.6$ if *flaming orb* (x) = *smouldering orb* (x) = 0.5, and for all other ordinary shades P , $P(x) = 0$
 $= 0.5$ otherwise.

(The difference between the definitions of Q^{**} and Q^* is that the first '0.6' in the definition of Q^{**} was a '1' in the definition of Q^* .) Now a degree-1 instance of *flaming orb* — an object which intuitively *is* determinate in respect of colour — is a degree 0.5 instance of Q^{**} , so Q^{**} should not count as a point property. However Q^{**} meets not only condition (c), but (c*) as well! Yet all is not lost, because Q^{**} does *not* meet condition (b) (which requires point properties to have degree-1 instances). For Q^{**} — or any other potential counterexample one might concoct — to satisfy (b), there would have to be a clause saying how something gets to be a degree-1 instance of Q^{**} . But given any such clause, Q^{**} will either no longer satisfy (c*) (if what it is to be a degree-1 instance of Q^{**} is, in whole or in part, to be a degree- r instance of some other point-like colour property); or it will no longer satisfy (a) (if what it is to be a degree-1 instance of Q^{**} is to be a degree- r instance of some *non*-point-like colour property such as *green*); or it will no longer be a *colour* property at all (if what it is to be a degree-1 instance of Q^{**} is to be a degree- r instance of some non-colour property such as the property of being *exactly one metre tall*).


Thus, in the final analysis, we adopt conditions (a), (b) and (c*) as our official account of what it takes to be a point property. Note that on this account, Tom is not

a degree-1 instance of *any* point property (for although Tom *is* a degree-1 instance of Q , and Q meets condition (c), Q does not meet condition (c*)). Tom is poised between two fully determinate shades of colour. Hence on this account we can simplify our definition of objective indeterminacy. We can say that x is indeterminate in respect of F just in case x is not a degree-1 instance of any point property in F .

THE WORLD AS A WHOLE

We have said what it is for an object to be indeterminate. It remains to say what it means for the world as a whole to be indeterminate. Perhaps the most natural proposal is to say that the world as a whole is indeterminate just in case it contains a genuinely indeterminate thing. But there are excellent reasons to resist this neat proposal.

Suppose that the property of weighing exactly 1.36 tons is a point property in the family of masses, and that the masses constitute a category. Now let C be a cloud and let it be genuinely indeterminate whether C weighs exactly 1.36 tons. Our account then implies — correctly, or so it seems — that C is a genuinely indeterminate entity. But should it follow from the fact that C is indeterminate in this way that the world as a whole is indeterminate in some respect? It would seem not. Our stipulations are consistent with a picture according to which the world is ultimately composed by a perfectly determinate array of democritean atoms, each sharply bounded and perfectly well defined as to position, momentum, colour, and whatever other fundamental physical parameters there may be. In a world of this sort, C gets to be indeterminate in mass by being indeterminate in composition. Suppose the atoms in W are arranged as in the picture:



World W

A is not quite close enough to the rest of *C* for it to be a determinate part of *C*, and not quite far enough away to be determinately distinct from it. And yet each atom in *W* is a perfectly determinate thing, standing in perfectly determinate relations to various other determinate things. And given this, it is natural to suppose that *W* ought to count as a fully determinate world.

W contains an indeterminate entity, but its indeterminacy emerges, as it were, from a more fundamental stratum of utterly determinate fact. Given that the atoms are arranged as they are in *W*, it is inevitable that it be indeterminate whether *A* is part of *C*, and hence whether *C* weighs exactly 1.36 tons.¹⁰ It seems natural to suppose that when an object's indeterminacy emerges in this way from an array of determinate facts and things upon which its composition and character supervene, its existence as an indeterminate thing does not compromise the determinacy of the world as a whole.

This may suggest that genuine worldly indeterminacy requires indeterminacy at the level of simples. But there are two reasons to resist this thought. In the first place, we should want to be able to entertain the question of worldly indeterminacy without presupposing that the world is ultimately constituted by an array of metaphysical atoms. But more importantly, consider a world *W** which contains exactly two atoms, *A* and *B* and nothing else. Both are determinate in all pertinent intrinsic respects — shape, mass, etc. But the distance between them is indeterminate in the following way. They're definitely not more than 1.1 mm apart, and they're definitely not less than 0.9 mm apart. But for each ordinary determinate distance relation *D* in between, $0 < D(A,B) < 1$. In particular, *A* and *B* are a 0.5 instance of the relation of being exactly one millimetre apart.

This strikes us as a fairly clear example of one way in which the world as a whole might be indeterminate. Of course it's strange. We should hope it would be. And yet the description is not obviously incoherent. But now ask, 'Does this world contain an intrinsically indeterminate atomic constituent?' The only atoms in the picture are *A* and *B*. (Assume we're talking about a Leibnizian world without substantial spacetime points or regions.) And for all we've said, the only respect in which either is indeterminate is its distance from the other. Since *being at such and*

such a distance from B is not an intrinsic property of A , we seem to have an indeterminate world whose constituent atoms are intrinsically determinate.

The examples suggest the following partial account. Suppose the world decomposes into an array S of simples, in the sense that every object in the world is ultimately composed of simples in S . These simples will have intrinsic natures, and they will stand in certain 'structural' external relations to one another. The world as a whole is then determinate if:

- (a) Every simple in S is a determinate object in the sense defined above, and
- (b) whenever some simples stand in the relevant external relations — relations of spatiotemporal distance, for example — there is no point relation in the relevant family of which they are an indeterminate instance.¹¹

This principle counts the cloud world, W , as a determinate world (despite the cloud), and it counts the two atom world W^* as an indeterminate world even though each atom is intrinsically determinate. So far, so good.

It remains to relax the assumption that the world decomposes into an array of simples. Here is one strategy. Say that Δ is a **division** of the world iff it is a collection of objects — not necessarily simple, not necessarily disjoint — such that every object in the world is ultimately composed of items in Δ . Now say that Δ is determinate iff it satisfies analogues of our two conditions:

- (a*) Every object in Δ is a determinate object in the sense defined above, and
- (b*) for any $x_1 \dots x_n$ in Δ , there is no point relation R in the relevant family of external relations such that $0 < R(x_1 \dots x_n) < 1$.

Our suggestion is that the world as a whole is determinate iff it has a determinate division.

¹⁰ This assumes that the principles of mereological composition are necessary truths. For some doubts about this, see [Rosen forthcoming].

¹¹ Given a family R of n -place relations, and an associated exact similarity relation \approx_R on n -tuples, the definition of a point-relation in R will then follow our earlier definition of a point property in F .

One final putative counterexample will motivate one final clarification of the proposal. Some philosophers believe in *seriously emergent properties* in the following sense. An intrinsic property *P* is seriously emergent just in case a complex thing may have it, and whether or not a complex thing has it is not settled by the natures of its constituent proper parts and their external relations to one another. Let's understand this as the denial of a supervenience claim. If *P* is seriously emergent, there might be two complex objects, *X* and *Y*, perfectly alike in all 'micro' respects (e.g., constituted by identical arrays of atoms) but which nonetheless differ as to *P*. We cannot rule out the existence of seriously emergent properties without some sort of argument. (Why should the intrinsic properties of the whole always be settled by the natures of its smaller parts and their relations to one another?) And since we do not know what such an argument would look like, we shall assume that they are possible, even if we can think of no credible examples.

Let's pretend that colour is seriously emergent in this sense. The world as a whole might then have a colour — the world might be a flaming orb — though not in virtue of the disposition of its parts. And if the world might have a colour, it might be indeterminate as to colour. It might be neither clearly flaming orb nor clearly not flaming orb. Surely a world of this sort ought to be reckoned an indeterminate world. But it's not obviously indeterminate given our account. It might well have a determinate division, or so it seems. Take the natural division into atoms, or if there is none, a division into very small bits, none of which is big enough to have a colour.

The natural response to this far-fetched possibility is to say that when seriously emergent properties are on the scene, the whole is (in one obvious sense) something over and above its parts. The nature of the whole is not settled by the natures of the parts and their external relations to one another. And when this is so, we can say that the whole is not *ultimately constituted* by the parts in the sense relevant to our principle. A genuine division is a division of the world into objects such that the nature of every object is fully settled by the natures of the objects in the division and their structuring external relations to one another. If some large object possesses an emergent property, then any suitable division must contain not only the parts of that object, but also the object itself. If the world as a whole manifests an emergent property, then any division of the world must include the world as a whole as a member. When the notion of a division is understood in this way, we can retain our formula: The world as a whole is determinate iff it has a determinate division.

CONCLUSION

The initial problem was to make sense of the notion that an object, or the world as a whole, might be intrinsically indeterminate in some respect. We have been guided throughout by the assumption that any decent account of worldly indeterminacy should preserve our sense that an indeterminate object would be a strange thing indeed. We hope we have succeeded at least in this.

Our paradigm for the indeterminate object is an object that possesses length, but no determinate length, or colour, but no determinate colour. We explain what this might mean by supplying a general definition of the notion of a point property in a category of properties. Our substantive analysis then comes to this: An object is indeterminate in respect of F iff it is an indeterminate instance of some point property in F . We then explain worldly indeterminacy in terms of this notion. We say that the world is indeterminate iff it lacks a determinate division into determinate objects.

We have helped ourselves to a number of assumptions along the way. In particular, we have assumed:

- (i) that there exist vague properties and relations, not just vague predicates;
- (ii) that properties and relations can be ranged into relatively maximal families or categories, and that for each such family F , there exists a relation of exact similarity in respect of F ;
- (iii) that properties and relations admit of real definition, so that given any relation R , we can raise the question, 'What is it for the X s to stand in R ?' 'In what does the fact that $R(x_1 \dots x_n)$ consist?' An answer to a question of this sort is an *account* of the relation R . The point of the idiom is to permit us to ask whether the definition of one property involves another, and so to identify certain relations of ontological involvement or dependence among properties.

We regard these notions as intelligible, though in each case we concede that further explication would be desirable. It would be nice to do without any one of them, especially the third. But we can see no way to do this, and perhaps this is an interesting result in itself. If you think you can make sense of the idea that the world might be vague or indeterminate, and if the only explication of this notion one can give relies on the notion of real definition, then that is some evidence that this notion must itself make some sort of sense.

REFERENCES

- Goodman, Nelson 1983. *Fact, Fiction and Forecast*, 4th edn, Cambridge MA: Harvard University Press.
- Islam, Amitavo 1996. The Abstract and the Concrete, typescript.
- Johnston, Mark 1987. Human Beings, *Journal of Philosophy* 84: 59–83
- Lewis, David 1986. *On the Plurality of Worlds*, Oxford: Blackwell.
- Rosen, Gideon forthcoming. The Limits of Contingency, in *Identity and Modality*, ed. F. MacBride and Crispin Wright, Oxford: Oxford University Press.
- Smith, Nicholas J.J. forthcoming. Vagueness and Blurry Sets, *Journal of Philosophical Logic*.