CRISPIN WRIGHT

ON THE COHERENCE OF VAGUE PREDICATES*

I. INTRODUCTION

Frege came to believe that a language containing vague predicates was essentially defective – that it was philosophically intolerable that predicates should occur for which it was not always determinate whether or not they could truly be ascribed to an object. Expressions of this conception are scattered throughout his writings. When, more seldom, he argues for it, it is on the ground that logical transformations may fail when applied to sentences containing expressions whose range of application has been only partially defined. It is not just a matter of the Law of Excluded Middle. Let F be a predicate defined only among, and universally applying to, individuals which are G. Then anything G is F. But the contrapositive fails: we cannot say that anything not-F is not-G, since the concept of having or lacking F has been fixed only for things which are G.[1]

Frege does not seem to have spared a thought for the idea that vague terms might require a *special* logic. The vagueness of ordinary language is seen rather as a flaw both needful and capable of remedy. This conception was endorsed by Russell[2] in his widely-despised introduction to the Tractatus. Ordinary language is always more or less vague; a logically perfect language, however, is not vague at all, and ordinary language is deemed to approach fulfilment of its function – that of having meaning – only in proportion as it approaches logical perfection, i.e. in inverse proportion to its vagueness.

Of course, we have long since abandoned the Frege-Russell view of the matter. We no longer see the vagueness of ordinary language as a defect. But we retain a second-order wraith of the Frege-Russell view in the notion that, even if the senses of many expressions in natural language are not exact, there is a precise semantical description for a given natural language, i.e., a theoretical model of the information assimilated in learning a first language or, equivalently, of the conceptual

Synthese **30** (1975) 325–365.

equipment in whose possession mastery of the language may be held to consist. Even if 'bald', say, is imprecise, this does not require any inexactitude in an *account* of its sense.

We suppose our use of language to be fundamentally *regular*; we picture the learning of language as the acquisition of grasp of a set of rules for the combination and application of expressions. The task of a philosophical theory of meaning is naturally interpreted as that of giving a systematic account of the contribution that constituents of a complex expression make to its content; we are concerned, that is to say, when we attempt such a theory, with the nature of the path from familiarity with the senses of the sub-sentential components of a new sentence to recognition of the sense of the whole. Such a theory will normally only concern itself with describing the contribution of a constituent *qua* expression of a certain logical type; it is in this connection that problems to do, e.g., with the nature of the distinction between proper names and other singular terms, or that between singular terms generally and predicative expressions, or with the question whether the notion of reference may properly be extended to apply to predicative expressions, derive their interest. The completion of such a theory would thus only be a preliminary to a full semantic description of a natural language; for we think not just of the type of contribution but of the *specific* contribution which a constituent makes as determined by semantic rules.

It is no obstacle to such a conception that we cannot in general informatively state such rules, i.e., provide a statement which could be used to explain the sense of an expression to someone previously unfamiliar with it. Consider, for example, the following schematic rule for a one-place predicate, F:

> 'F may truly be applied to an individual, a, if and only if a satisfies the condition of being $\emptyset$'.

How should we specify $\emptyset$ if F is 'red'? Plainly the only such condition which actually captures our understanding of 'red', rather than, in a merely extensional way, its conditions of application, is ... the condition of being red. So we may not in general expect instances of the schematic rule to be of any explanatory use if they are stated in a given language for a predicate of the same language; as a corollary, it will not in general be possible to appeal to such a rule to settle a dispute about the ap-

plicability of an expression. It remains open to us, nevertheless, to regard such a rule as exactly encapsulating (part of) what is understood by someone who understands, e.g. 'red', for it states conditions recognition of whose actualisation is sufficient to justify him in describing an object as 'red'; it is merely that such a capacity of recognition cannot be bestowed by stating the rule.

Our picture, then, is that to use language correctly is essentially nothing other than to use it in conformity with a set of instructions, of semantic rules. Of course our handling of language is in general quite automatic, but so is a chess player's recognition of the moves allowable for a piece in a certain position; it remains true that an account of his knowledge is to be given by reference to the rules of chess.

The question now arrises, what means are legitimate in the attempt to discover features of the *substantial* rules for expressions in our language, the rules which determine specifically the senses of such expressions? The view of the matter with which we are centrally concerned in this paper is that we may legitimately approach our use of language from within, that is, reflectively as self-conscious masters of it, rather than externally, equipped only with behavioural notions. We may appeal to our conception of what *justifies* the application of a particular expression; we may appeal to our conception of what we should count as an adequate explanation of the sense of a particular expression; to the limitations imposed by our senses and memories on the kind of instruction which we can actually implement; and to the kind of consequence which we attach to the application of a given predicate, to what we conceive as the point of the classification which the predicate effects. The notion that forms the primary concern of this paper – henceforward referred to as the *governing view* – is that we can derive from such considerations a reflective awareness of how expressions in our language are understood, and so of the nature of the rules which determine their correct use.

The governing view, then, is a conjunction of two theses: that our use of language is properly seen, like a game, as an activity in which the allowability of a move is determined by rule, and that properties of the rules may be discovered by means of the sorts of consideration just described. The governing view does not involve a psychologistic conception of understanding, according to which understanding would be regarded as an essentially mental state of which the correct employment

of an expression was a mere behavioural symptom; but it provides means in excess of the behaviourist minimum, = the description of when an expression is actually used, for an attempt to investigate the nature of the semantic rules operative with respect to a given class of expressions. It may be that, as so far characterised, the view will seem platitudinous; the purpose of this paper is to question its coherence.

A striking feature of the *Philosophical Investigations* is the hostility displayed by the author of the notion of a language-game to the idea that it is *explanatory* of our use of language to appeal to the concept of rule-following. It is not, of course, that there cannot be a resoluble dispute about the correct use of a particular expression. Rather Wittgenstein seems to argue for an indeterminacy in the *identity* of the rules which someone supposedly follows. Here it is irrelevant whether we can supply, e.g., what seems an informative completion of the above schematic rule for a particular predicate, *F*. The question is, what content is there to the claim that such a rule faithfully incorporates someone's understanding of *F*? For he and we may sincerely agree on a particular formulation and then, sincerely and irresolubly, apply *F* in mutually inconsistent ways. It is useless to protest that all that follows is that he uses *F* in accordance with a different rule. We agreed on how the rule was to be characterised; now, it seems, we reserve the right to offer some other characterisation of what governs *his* use of the expression. (Though there is no necessity that any such characterisation should occur to us.) From his point of view, however, the initial characterisation remains perfectly adequate; it is of *our* use of *F* that a re-characterisation of the determining rule is required. But, now, what objectivity is there in the idea of the *correct* characterisation of his or our respective rules? And, if there is none, how can we penetrate, as it were, to the real nature of the rule which *we* follow, if we wish to explain our use of *F*? Thus the terminology of semantic rules, adequate perhaps for certain purposes as a picture, fails to express a domain of objective fact from which our use of language may be seen to flow.

This line of argument poses a difficulty for the first thesis of the governing view. It is mentioned as a familiar type of criticism of the notion of a semantic rule, and by way of contrast with the character of the difficulty for the governing view to be posed in this paper. Here it is contended rather that the second thesis, concerning the means whereby

features may be discovered of the semantic rules which we actually follow, constrains us to recognise *semantic incoherence* in our understanding of a whole class of predicates – elements whose full exploitation would force the application of these expressions to situations where we should otherwise regard them as not applying. The second thesis requires the recognition of rules which, when considered in conjunction with certain general features of the situations among which their associated expressions are to be applied, issue in contradictory instructions. Nevertheless we succeed in using these expressions informatively, and to use language informatively depends on using it, in large measure, consistently. It follows that our use of these expressions cannot correctly be pictured purely as the implementation of the rules which the second thesis yields for them – these rules cannot be obeyed by *consistent* behaviour. The governing view is therefore incoherent; for if its second thesis is true, the semantic rules which are operative with respect to certain predicates are capable by consistent beings only of selective implementation and thus, contrary to the first thesis, are not *constitutive* of what we count as the correct use of these expressions.

Predicates of the relevant kind are all examples of a certain sort of vagueness: not exactly borderline-case vagueness, if that is understood, as Frege sometimes describes it, as the existence of situations to which it is indeterminate whether or not a predicate applies, but something which, under the guise of a favourite metaphor, he constantly runs together with the possession of borderline-cases, viz., the idea of lacking 'sharp boundaries', of dividing logical space as a blurred shadow divides the background on which it is reflected. The conflation is plausible because the image equally exemplifies the idea of the borderline-case, a region falling neither in light nor in shadow. But there seems no reason why having borderline-cases should imply blurred boundaries. Borrowing another of Frege's analogies, we may assimilate a predicate to a *function* taking objects as arguments and yielding the *True* or the *False* as values; in these terms a predicate with borderline-cases is simply a partial such function, and that is consistent with the obtaining of a perfectly sharp distinction between cases for which it is defined and cases for which it is not. Borderline-case vagueness *simpliciter* presents no difficulty for the governing view; it is merely that we are presented with situations to which no response is determined by the semantic rules of our language

as the correct one. On the other hand, if the second thesis of the governing view is correct, then predicates with 'blurred boundaries' are, in typical cases, to be regarded as semantically incoherent.

This incoherence resides in their vagueness *as such*. It is plausible to suppose that the vagueness of many expressions is not, as Frege and Russell thought, merely a reflection of our intellectual laziness. Rather, the utility and point of the classifications expressed by many vague predicates would be frustrated if we supplied them with sharp boundaries. (If it is an empirical truth that stress diseases are more widespread in highly concentrated populations, it is doubtful whether it would survive an exact numerical definition of 'highly concentrated'.) The sorts of consideration admitted by the second thesis will transpire, in the succeeding sections of the paper, to yield support for the idea that such predicates are *essentially* vague. The thesis equips us to argue that lack of sharp boundaries is not just a surface phenomenon reflecting a hiatus in some underlying set of semantic rules. Lack of sharp boundaries is not the reflection of an omission; it is a product of the kind of task to which an expression is put, the kind of consequences which we attach to its application or, more deeply, the continuity of a world which we wish to describe in purely observational terms. Lack of sharp boundaries is semantically a deep phenomenon. It is not generally a matter simply of lacking an instruction where to draw the line; rather the instructions we already have determine that the line is *not* to be drawn. This conclusion might seem a welcome contribution to our understanding of the nature of vagueness, were it not that it is supplied in the form not merely that no sharp distinction may be drawn between cases where such a predicate applies and cases where it does not, but that no such distinction may be drawn between cases where it is definitely correct to apply the predicate and cases of *any* other sort.

Thus it is that an adherent of the governing view simply has no coherent approach to the Frege-Russell view of vagueness. His second thesis requires him to reject the suggestion that vagueness is a superficial, eliminable aspect of natural language with no real impact upon its informative use. But it does so by means of considerations which require him to regard many vague predicates as semantically incoherent, so that, unless the Frege-Russell view is right, he cannot maintain his first thesis with respect to such expressions. Only if their vagueness *is* an incidental

feature can he maintain that the *essential* semantics of such expressions conform to his first thesis.

The programme for the remaining part of the paper is as follows. In the next section, three examples will be presented of predicates to which the second thesis attributes semantic incoherence; specifically, it sustains for each of them the reasoning of the Sorites paradox. The character of the incoherence will be generalised, and arguments afforded by the second thesis for such an account of the semantics of these expressions will be developed. In Section III a fourth and deeper rooted example will be presented. In Section IV we shall consider an obvious strategy for solving paradoxes in the Sorites group which, if allowed, would undercut the considerations of Sections II and III. This strategy will be rejected. In Section V it will be argued that one class of predicate to which the second thesis attributes semantic incoherence is, in a certain sense, ineliminable. It will be shown that a seemingly promising adaptation of certain of Goodman's[3] ideas fails to provide an adequate refashioning of the semantics of these predicates; and that a simpler suggestion, while indeed liberating them from semantic incoherence, does so at the cost of generating other predicates with the same feature. In the final section our conclusions will be drawn together.

If the governing view is unacceptable, that is something which it is as well to know. The interest of the paper, however – if any – derives equally from that of the issues which we shall have occasion to discuss *passim*: the nature of vagueness and the correct logic for vague expressions; ostensive definition and observational language; the empirical sources of the concept of continuity, and the notion of order within phenomenal continua like the spectrum of colours.

II

Example 1

Our first example is the Megarian paradox itself.[4] If we begin with a pile of salt large enough to be fairly described as a heap, the subtraction of a single grain of salt cannot make a relevant difference; if $n+1$ grains of salt may constitute a heap, so may n such grains. Of course, the plausibility of this supposition derives from the species of sense with which 'heap' has been endowed; we have not fixed exact boundaries for the concept of a heap, either in terms of the precise number of grains con-

tained or, indeed, in terms of any other precise measure. To allow that at some stage the subtraction of a single grain might transform a heap into a non-heap would be exactly to anticipate the determination of such a specific boundary.

However, in the present semantic situation of 'heap' it would be merely an error to insist that, for some particular value of n, $n+1$ grains of salt would amount to a heap while n grains would not; for that is simply not the sense of 'heap'. If it were agreed in some particular case that $n+1$ grains did amount to a heap, no-one could produce a telling reason for withholding the predicate from the same pile minus one grain; except perhaps avoidance of the incoherence implicit in the situation.

'Heap', then, would appear to be semantically incoherent, for its sense is such that it essentially lacks exact extensional boundaries; and its lack of such boundaries demands particularly that a transition from $n+1$ grains to n grains can never be recognised to transform a case where 'heap' applies to a case where it does not. So we are seemingly equipped to force the application of 'heap' through successively smaller aggregations of salt-grains, terminating in items which amount not to heaps but merely, say, to pinches. Here we gravitate towards the idea that lack of exact boundaries is, as such, an essentially incoherent semantic feature. This idea, however, will need qualification.

Example 2

Predicates of degree of human maturity – 'infant', 'child', 'adolescent', 'adult', display the same peculiarity. They are mutually inconsistent yet lack sharp boundaries with respect to their neighbours in the scale of human maturation. More exactly, if we take some sufficiently small interval of time and suppose that someone matures in a typical fashion, then at no stage will he effect *within* such an interval of time a transition from one stage of maturity to the next.

To illustrate the point with an example from Esenin-Volpin:[5] take as the relevant interval the span of time from one heartbeat to its successor; then the concept of childhood – the sense of 'child' – is such that one does not, within a single hearbeat, pass from childhood to adolescence. To be sure, one is not a child forever; but at least childhood does not evaporate between one pulse and the next. If one's nth heartbeat takes place in childhood, then so does the $n+1$th[6].

'Infant', 'child', 'adolescent', 'adult', are thus all semantically incoherent expressions; for the sense of each of these predicates is such that, in a typical process of growing-up, their correct application will always survive the transition from one heartbeat to its successor or to its predecessor. So again, by appropriately many steps of *modus ponens*, we may force the application of each of these predicates to cases we should otherwise regard as falling within the domain of a competitor.

Example 3

Consider a series of homogeneously coloured patches, ranging from a first, red patch to a final, orange one, such that each patch is *just* discriminable in colour from those immediately adjacent to it, and is more similar to its immediate neighbours than to any other patches in the series. Marginal, uni-directional changes of shade are thus involved in every transition from a patch to its successor. Now, it is notable that the sense of colour predicates is such that their application always survives a very small change in shade. Given that one is content to call something 'red', one will still be so content if its colour changes by some just discriminable amount. There is a notion of a degree of change in respect of colour too small to amount to a change *of* colour. Only if a substantial difference intervenes between two patches shall we consider ourselves justified in ascribing to them incompatible colour predicates.

Obviously this is an incoherent notion. In particular in view of the proximity in shade of neighbouring patches in our series, it provides an easy proof that all the patches are red, (or orange, or doubtfully either). Moreover any two colours can be linked by such a series of samples; so any colour predicate can likewise be exported into the domain of application of one of its rivals. Colour predicates as a class are semantically incoherent.

In these examples we encounter the feature of a certain *tolerance* in the concepts respectively involved, a notion of a degree of change too small to make any difference, as it were. There are degrees of change in point of size, maturity and colour which are insufficient to alter the justice with which some specific predicate of size, maturity or colour is applied. This is quite palpably an incoherent feature since, granted that any case to which such a predicate applies may be linked by a series of 'sufficiently

small' changes with a case where it does not, it is inconsistent with the exclusivity of the predicate.

More exactly, let $\emptyset$ be a concept related to a predicate, F, in the following way: that any case to which F applies may be transformed into a case where it does not apply simply by sufficient change in respect of $\emptyset$; colour, for example, is such a concept for 'red', size for 'heap', degree of maturity for 'child', number of hairs for 'bald', degree of complexity for 'memorable' as applied to patterns, and so on. Then F is tolerant with respect to $\emptyset$ if there is also some positive degree of change in respect of $\emptyset$ insufficient ever to affect the justice with which F is applied to a particular case. These wholly intuitive ideas are sufficiently clear for present purposes.

At this point it has to be conceded that the manner in which these examples have been presented has been wholly tendentious. It is not to be doubted that the predicates in question do lack sharp boundaries; and the antiquity of the paradox bears witness to how easy it is to interpret this as involving the possession by these predicates of a principle of re-application through marginal change. But is this a correct interpretation? If 'heap', for example, lacks sharp boundaries, then certainly we are not equipped to single out any particular transition from n to $n-1$ grains of salt as being the decisive step in changing a heap into a non-heap; no one such step is decisive. But that is not to say that such a step always *preserves* application of the predicate. Should we not instead picture the situation as comparable to that in which neighbouring states fail to agree upon a common frontier? Their failure to reach agreement does not vindicate the notion that e.g. a single pace in the direction of the other country always keeps one in the original country. For they have at least agreed that there is to be a border, that *some* such step is to be decisive; what they have not agreed is where.

The analogue of this conception for the predicates which interest us is exactly that their vagueness is purely a reflection of our intellectual laziness. We have, as it were, decided that a disjunction is to be true – at some stage, n grains will be a heap whereas $n-1$ grains will not – without following up with a decision about *which* disjunct is true. The notion that these predicates are tolerant confuses a lack of instruction to count it the case that a proposition is false with the presence of an instruction to count it true. This conflation would be permissible only if our set of semantic rules were in a certain sense complete, that is, if it contained instructions

for every conceivable situation. But for there to be vague expressions in our language is precisely for this not to be so.

If we avail ourselves of the types of consideration afforded by the governing view, we shall reject this suggestion as a deep misapprehension of the nature of this species of vagueness. The lack of sharp boundaries possessed by these examples is correctly interpreted as tolerance, provided that we may so discover elements of their senses.

The point is not that there is any decisive formal obstacle to providing general stipulations by means of which we might identify in such examples a last point at which the predicate in question could definitely correctly be applied. We could set upper and lower limits for 'heap' in terms of number of grains.[7] We could do the same for 'child', etc., by setting precise age limits to the successive phases of life – perhaps in terms of numbers of heartbeats. For colour predicates it would *prima facie* be less easy to provide such a refinement, since the notion of being of the same shade encounters difficulties which those of being of the same age, or containing the same number of grains, do not. But suppose it is possible. Then what in the semantics of these examples is already inconsistent with our so refining their senses?

'Heap' is essentially a coarse predicate, whose application is a matter of rough and ready judgement. We should have no use for a precisely demarcated analogue in the contexts in which the word is typically used. It would for example be absurd to force the question of the execution of the command, 'Pour out a heap of sand here', to turn on a count of the grains. Our conception of the conditions which justify calling something a heap of sand is such that the justice of the description will be unaffected by any change which cannot be detected by *casual* observation.

A different argument is available for colour predicates. *We* learn our basic[8] colour vocabulary ostensively, that is, by exposure to samples illustrative of its application. Evidently it is a precondition of our capacity to do so that we can reasonably accurately remember how things look. We can imagine people who can recognise which simultaneously presented objects match in colour, and so are able to use a colour-chart, but who cannot in general remember shades of colour sufficiently well to be able to employ without a chart a colour vocabulary anything like as rich as ours. Such people, for example, having been shown something yellow, might later be quite unable to judge whether what they were earlier

shown would match the orange sample now before them. Thus, for such a community, an ostensive definition of 'yellow' would not be feasible, and, in order to make the distinctions of our basic colour vocabulary, they would require to employ charts.

Of course, we also use charts for some purposes; but never to make distinctions of the magnitude of those, say, between the colours of the rainbow. Any object to which one of these predicates definitely correctly applies may be recognised as such just on the basis of our ostensive training. But if this is so, it has to be a feature of the senses thereby bestowed upon these predicates that changes too slight for us to remember – that is, a change such that exposure to an object both before the change is undergone and afterwards leaves one uncertain whether the object *has* changed, because one cannot remember sufficiently accurately how it was before – never transform a case to which such a predicate applies into one where such is not definitely correctly the right description. The character of the basic colour training which we receive, and which we hand on to our children, presupposes the *total memorability* of the distinctions expressed by our basic colour predicates; only if single, unmemorable changes of shade never affect the justice of a particular, basic colour description, can the senses of these predicates be explained entirely by methods reliant upon our capacity to remember how things look.

With respect to 'child', etc., the governing view allows us a third type of consideration. Broadly, we do not make distinctions in phases of maturity in the fashion of a naturalist, just to record a variety with which we are confronted; rather, these classifications are of substantial social importance in terms of what we may appropriately expect from, and of, persons who exemplify them. Thus infants have rights but not duties, whereas of a child outside infancy we demand at least a rudimentary moral sense; we explain the anti-social behaviour of some adolescents in terms of their being adolescents; and we make moral and other demands of character on adults which we would not impose on the immature.

It is plausible that the predicates of Example 2 could not endure this treatment, were they not tolerant with respect to marginal changes in respect of degree of maturity. It would be irrational and unfair to base substantial distinctions of right and duty on marginal – or even non-

existent[9] – such differences; if we are forced to do so, for example with electoral qualifications, it is with a sense of artificiality and absurdity.

There is another point: without tolerance these predicates could no longer sustain the explanatory role which they now have for us. Only if a *substantial* change is involved in the transition from childhood to adolescence can we appeal to this transition to explain substantial alterations in patterns of behaviour; if some adolescents differ barely, if at all, from some children, so that no significant change need be undergone in making the transition, to have done so can explain nothing. That predicates of degree of maturity possess tolerance is a direct consequence of their social role; very small differences cannot be permitted to generate doubt about their application without correspondingly coming to be associated with a burden of moral and explanatory distinctions which they are too slight to convey.

Our embarrassment about where to 'draw the line' with these examples is thus a reflection not of any hiatus in our semantic 'programme' but of the tolerance of the predicates in question. If casual observation alone is to determine whether some predicate applies, then items not distinguished by casual observation must receive the same verdict.[10] So single changes too slight to be detected by casual observation cannot be permitted to generate doubt about the application of such a predicate. Likewise, if the conditions under which a predicate applies are to be generally memorable, it cannot be unseated by single changes too slight to be remembered. Finally, very slight changes cannot be permitted to generate doubt about the application of predicates of maturity without contravening their moral and explanatory role.

The utility of 'heap', the memorability of the conditions under which something is red, the point of 'child' thus appear to impose upon the semantics of these predicates tolerance with respect to marginal change in their various relevant respects. Not that, on the governing view, those considerations provide a whole account of their vagueness. For example, the considerations applied to 'heap' presumably apply to 'red' and, in some measure, to 'child' also. However, it is clear that to allow just the foregoing sketchy considerations is to concede both that the vagueness of our examples is a phenomenon of semantic depth – that is, it is sacrificed at much more than the cost of the intellectual labour of the stipulation – and that it is a structurally incoherent feature. Two things

follow. First, there is no special logic for predicates of this sort, crystallising what is distinctive in their semantics in contrast with those of exact predicates; for what so distinguishes them is their inconsistency. A 'logic' of this species of vagueness is chimerical. Second, the manner in which we typically use these expressions needs some other model than the simple following of rules, if these are to incorporate all the features of their senses which we should wish to recognise. It is perhaps more nearly comparable to the behaviour of a public-spirited citizen in relation to fiscal law; for, as is familiar, the overall effectiveness of a system of taxation may well depend both upon the presence of loopholes and on people's forbearance to exploit them.

III

Anyone who holds the second thesis of the governing view should recognise a distinct and more profound source of tolerance in adjectives of colour than the inability of our memories to match the sharpness of discrimination possessed by our senses. Colour predicates, it is plausible to suppose, are in the following sense purely *observational:* if one can tell at all what colour something is, one can tell just by looking at it. The look of an object decides its colour, as the feel of an object decides its texture or the sound of a note its pitch. The information of one or more senses is decisive of the applicability of an observational concept; so a distinction exemplified in a pair of sensorily-equivalent items cannot be expressed by means solely of predicates of observation, for any observational expression applying to either item must apply to both. What is about to be illustrated with respect to colour-predicates is, under appropriate assumptions, a feature of *any* expression whose sense is observational in the fashion just sketched.

Since colour predicates are observational, any pair of objects indistinguishable in point of colour must satisfy the condition that any basic colour predicate applicable to either is applicable to both. It is, however, familiar that we may construct a series of suitable, homogeneously coloured patches, in such a way as to give the impression of a smooth transition from red to orange, where each patch is *indiscriminable* in colour from those immediately next to it; it is the non-transitivity of indiscriminability which generates this possibility. So, since precise matching is to be sufficient for sameness of colour, we can force the

application of 'red' to all the patches in the series, some of which are not red but orange. That is: since 'red' is observational, its sense must be such that from the premises, that x is red and that x looks just like y, it follows that y is red, no matter what objects x and y may be. This rule enables us to conclude that each successive patch in our series is red, given only the true premise that the first patch is red.

The purpose of Example 3 was to illustrate, from the standpoint of the second thesis, a tolerance of predicates of colour with respect to marginal changes of shade. If we retain such an account for this new example – Example 4 – we shall be forced to regard identity of shade as a non-observational notion. We shall be admitting that changes in shade take place between adjacent patches where none *seem* to have taken place, where the most minute mutual comparison reveals no difference. But it is clear that we shall be driven to some such admission even if, rather than construe it as non-observational, we abandon the notion of identity of shade altogether. For some sort of non-observational changes are taking place in this example, however we choose to describe them.

Example 4, if allowed, reveals colour predicates as tolerant with respect to changes which cannot be directly discerned in objects which undergo them; an object may suffer such a change without it being possible to discover that it has done so save by comparing it with something else. This feature differentiates this example from the others. Moreover we do not have ready to hand a concept in terms of which we can describe what these changes essentially are. Not that we could not offer an account in terms e.g. of the physics of light; but it is contingent that the changes in question are associated with any particular physical changes which we might independently discover. On the other hand it is seemingly not a full account of the matter just to say that a patch now matches something which it did not match before, while not itself seeming to have changed in the meantime; this is unsatisfactory as a full account because we are inclined to say that, when objects come into a relation which they did not share before, one or both of them must have changed in independently specifiable respects; whereas, in this case, *neither* patch may seem to have changed. We are lacking, for example, a notion standing to the concept of shade as that of real position stands to phenomenal position. Had we such a concept, we could account for the changes in question in real terms.

This lack would be a *prima facie* obstacle to stipulating away examples of this fourth type; for the stipulations envisaged in the other three cases each made use of a concept – number of grains, number of heartbeats, phenomenal shade – in terms of which we could describe the small changes involved. Later we shall consider how this obstacle might be overcome. For the present it is more significant that the difficulty arises at all. Had we a notion of shade comparable to that of real position, we might provide such a stipulation; but had we such a notion of shade, we should already have surrendered *its* observationality, and to provide the stipulation would be to surrender that of colour as well. To stipulate away the tolerance of a predicate is to provide a general explanation of where, in a series of the relevant sort, it may be applied correctly for the last time. But if we did so with respect to 'red', 'orange' and the type of series illustrated in Example 4, we should have to forego our entire present conception, viz. the *look* of a thing, of what justifies the application of these predicates. To say that a predicate is observational requires that the conditions under which it may be applied to an object must be determinable simply by observation of it; whereas we are contemplating a situation where 'red' may definitely correctly be applied to only one of a pair of cases which are observationally exactly similar – whose looks match exactly. Clearly, then, any observational predicate must display tolerance in a series in which it is initially but not everywhere exemplified and whose every member cannot be observationally distinguished (save numerically) from those immediately adjacent to it. Naturally, no series can satisfy these conditions unless the relevant relation of indistinguishability behaves non-transitively.

These considerations are broadly analogous to what was said of the Heap: if we so fix the sense of a predicate that whether it applies has to do with nothing other than how an object seems when casually observed, then changes other than can be determined by casual observation of it cannot transform a case to which the predicate applies into one where it does not, or to which its application is somehow doubtful. The point remains good if we omit the word, 'casual'. But Example 4 is seemingly deeper-reaching, at any rate for someone who, in the spirit of Frege, required that language should be purified of vague expressions. The cost of eliminating predicates of casual observation is convenience; but the cost of eliminating the use of expressions tolerant in the manner of

Example 4 would be the abandonment of predicates of strictly observational sense. Might there not then be a higher price to pay, namely, the jeopardising of contact between language and empirical reality?

We shall not pursue this thought immediately. First we need a deeper perspective for Example 4. We require an explanation of the observationality of certain predicates – or, what in this context comes to the same thing, some reinforcement of the supposition that their semantics *are* purely observational – and we require to know under what circumstances we may expect our sensory discriminations to be non-transitive. The intersection of these explanations will yield a general indication of the range of the type of case which Example 4 illustrates.

That we do intuitively regard the semantics of adjectives of colour as purely observational is beyond doubt; and simply illustrated by the fact that we should regard it as a criterion of lack of understanding of such an adjective if someone was doubtful whether both of a pair of objects which he could not tell apart should receive the same description in terms of it. We regard it as a criterion of understanding such a predicate that someone, presented under suitable conditions with an object to which it applies, can tell that it does so just on the basis of the object's appearance. But can an explanation be provided of why there should be predicates with such a semantics?

It might be supposed that any *ostensively definable* predicate must be observational. If an expression can be ostensively defined, it must be possible to draw to someone's attention those features in his experience which warrant its application; and if this is possible there can be no question of the expression applying to some but not others among situations which he cannot distinguish experientially. It would be a poor joke on the recipient of an ostensive definition if the defined expression applied selectively among situations indistinguishable from one which was originally displayed to him as a paradigm.

Unless it is disallowed that aspects of the semantics of an expression can be discovered by appeal to such considerations, we are bound to regard this suggestion as basically correct. Accordingly, we can no longer combine the conventional idea of the place which ostensively defined predicates occupy in our 'conceptual scheme' – the base of the epistemic pyramid – with the view that strict semantic coherence is a necessary condition of intelligibility. Rather, we have to recognise that such

predicates are endowed by their very mode of introduction with a kind of original sin – a species-liability to tolerance.

We are not yet, however, in a position to draw this conclusion. Certainly an ostensive definition must be regarded as issuing a license, so to speak, to apply an expression to any situation relevantly matching that which the definition uses; an ostensive definition must do this or it is of no use. But that is not to issue a license to apply the expression to a situation which does *not* match the original situation, but merely matches an intermediate situation, indistinguishable from both. If we think of an ostensive definition as a command: 'Apply F to situations like this', the command obviously does not apply to situations which are not 'like this' but merely like something like this. This is so whether the required resemblance is conceived as indistinguishability or as some less exact likeness.

What, then, is the connection between an expression's being observational – its applying to both, if to either, of any pair of observationally indistinguishable situations – and its being ostensively definable? It is as follows. The picture of acquiring concepts by experience of various cases where they do apply and various cases where they do not – a picture which surely has *some* part to play in a philosophically adequate conception of the learning of a first language – cannot be wholly adequate for concepts which differentiate among situations which look, feel, taste, sound and smell exactly alike. So if that picture is wholly adequate for any concepts, they must be concepts whose range of application does not include situations which experience cannot distinguish from situations which may not definitely correctly be regarded as falling within that range. To master the sense of a predicate is, at least, to learn to differentiate cases to which it is right to apply it from cases of any other sort. If such mastery can be bestowed ostensively, a comparison of two such cases must always reveal a difference which sense-experience can detect. The notion, then, that the whole range of application of a predicate can be made intelligible by ostensive means, presupposes that it is never the case that only one of a pair of objects, which the senses cannot tell apart, is characterised by it.

It is tempting to suggest that the trouble ultimately resides in our whole conception of what an ostensive definition tries to achieve. The point of an ostensive definition of 'red' is to communicate the concept

of a certain look, a look which is to be understood as invariably justifying calling something which has it, 'red'. Looks, however, like sounds and smells are phenomenal, so that they are imbued with the same instability which was noted earlier in the idea of phenomenal position. If it is sufficient to share a look that things seem exactly similar, then the non-transitivity of indiscriminability provides a way of proving that everything has the same look. If it is not so sufficient, then the nature of a particular look cannot be revealed by mere display of something which has it, for something could appear absolutely similar yet not have the look; the 'look' of something is no longer a purely phenomenal notion, so it cannot be communicated by ostensive procedures, for all that they can capitalise upon is how things appear. One cannot, for example, give an ostensive definition of a real length.

Be that as it may, there is a clear, general connection between observationality and ostensive definition. If there is in the conditions of correct application of a predicate nothing which is incapable of ostensive communication, then the predicate must apply to both, if to either, of any pair of indistinguishable objects. But it seems manifest that adjectives of colour, and many others, do precisely not involve any such further condition of correct application; on the contrary, ostensive training would appear fully determinant of their meaning – or, if it is not, it is the only training which we get.

Let us move, then, to the second question: under what circumstances can we expect our sensory discriminations to be non-transitive?

It is natural to view it as a consequence of the *coarseness* of our perceptions that series of colour samples can be constructed in which we can directly discern differences in hue only between non-adjacent members. For we do admit in many cases, e.g. the concept of spatial position, the idea of a change too small to be directly perceived, *par excellence* too small to be perceived without special apparatus. This admission entails that indiscriminability will behave non-transitively in suitable circumstances, since small, imperceptible differences may add up to a larger, noticeable one.

But it might not seem that this explanation is the one which we seek; for it presupposes the admission of a species of *objectivity* which cannot be exemplified by observational predicates. The observationality of colour words requires that looking the same colour is sufficient for the same

colour predicate to apply: whereas the proposed explanation, applied in the case of spatial position, essentially distinguishes between when things *seem* to be in the same position and sufficient conditions for their being truly described as being so. A change of position need not be noticeable; a change of colour must be. Naturally we may be driven to abandon this very aspect of the semantics of colour predicates as a result of the tension between observationality and non-transitive indiscriminability; but until we do so, it cannot be an *explanation* of the non-transitive indiscriminability of suitable colour samples that differences of colour exist too fine for our gross perceptions.

This however is still not quite the point. For plainly, as we noted, the non-transitivity of matching requires that not *every* feature of colour patches can be a directly observational one; colour patches evidently allow of changes, whether these changes are described as changes *in* colour or not, of a kind which cannot be directly discerned. So we have no alternative but to admit a gap among such items between seeming not to have changed in any respect, and actually not having done so. We have to admit the described objectivity, if not in the notion of change in colour, at least in the notion of change in some respect. What is really wrong with the explanation is not the presupposition of objectivity, but the circumstance that the sole ground for affirming that there is a distinction between seeming and being here is that indiscriminability is behaving non-transitively; whereas we are trying *independently* to circumscribe the circumstances under which we may expect such behaviour. So we require an alternative account.

Suppose that we are to construct a series of colour patches, ranging from red through to orange, among which indiscriminability is to behave transitively. We are given a supply of appropriate patches from which to make selections, an initial red patch C_1, and the instruction that each successive patch must either match its predecessor or be more like it than is any other patch not matching it which we later use. Under these conditions it is plain that we cannot generate any change in colour by selecting successive matching patches; since indiscriminability is to be transitive, it will follow if each C_i in the first n selections matches its predecessor, that C_n matches C_1. The only way to generate a change in colour will be to select a non-matching patch.

When $\langle C_1, \ldots, C_n \rangle$ is complete, how will it look in comparison with

the series of Example 4? It is clear that we shall have lost what is distinctive of that series: the appearance of *continuous* change from red to orange. In the new series the shades are exemplified in discrete bands (containing perhaps no more than one patch) and all the changes take place abruptly in a transition from a patch to its successor. So it appears that, were our judgements of indiscriminability to be universally transitive among samples of homogeneous colour, no field of colour patches could be ordered in the distinctive fashion now possible; that is, so as to give the impression of a perfectly smooth change of hue. If matching generally behaved transitively among shades, no series of colour patches could give the impression of continuous transformation of colour; contrapositively, then, for matching to function non-transitively among a finite set of colour patches, it is sufficient that they may be arranged so as to strike us as forming a phenomenal continuum.

This reasoning may obviously be generalised. Any finite series of objects, none of which involves any apparent change in respect of $\emptyset$, may give an overall impression of continuous change in respect of $\emptyset$ only if indiscriminability functions non-transitively among its members.

Evidently, though, not all processes of seemingly-continuous change come, as it were, ready made out of finitely many stages. Seemingly-continuous processes in time do not generally do so; and nor do certain purely spatial seemingly-continuous changes, e.g. the convergence of a pair of near-parallel lines. The question now arises whether there is not a more general connection between seemingly-continuous change and non-transitive indiscriminability than that illustrated in the somewhat artificial case of the colour-patches.

Let us consider the case of processes of change in time. Let D be such a process, and let a *stage* of D be the state of D at a particular point in time, an instantaneous exposure, as it were, of the process at that point. D is to be *non-recurrent* in the following sense: if $\emptyset$ is the respect in which changes in D take place, no distinct stages, x, y, z in D are to be such that z is in respect of $\emptyset$ more like x than y is when $x<y<z$; ('$<$'=is earlier than). It is analytic that any process of change is non-recurrent up to some stage. We claim that, on a very natural presupposition, D will be seemingly continuous only if there is some finite selection of stages of D among which indiscriminability behaves non-transitively.

Suppose on the contrary that while D itself is seemingly-continuous,

there is no way of selecting stages from D, however close together, so that indiscriminability behaves non-transitively among them. Consider a maximal set, S, of stages of D yielded by the following rule of selection: x is a member of S if and only if x is discriminable in respect of $\emptyset$ from every y previously selected as a member. (Let the first selection be regarded as satisfying this condition vacuously.) Since S is maximal, every stage in D must be indiscriminable in respect of $\emptyset$ from some member of S. Suppose now that S is finite and consider the series, $\langle D_1, \ldots, D_n \rangle$, of all its members in order of temporal succession. Let D_i, D_{i+1}, be a pair of stages adjacent in this series.

Plainly any stage occurring in D later than D_i, but earlier than D_{i+1} must be indiscriminable in respect of $\emptyset$ from either D_i or D_{i+1} but not both; it cannot match both, since matching is hypothesised to be behaving transitively; it cannot match neither, or it must match some other stage in $\langle D_1, \ldots, D_n \rangle$, so violating the hypothesis that D is non-recurrent. Clearly if such an in-between stage matches, say, D_i, then all stages lying temporally between it and D_i in D must likewise match D_i, or D will not in this region be non-recurrent; *mutatis mutandis* if it matches D_{i+1}. So the region of D between D_i and D_{i+1} must divide into two contiguous segments, every stage in one of which will match D_i while every stage in the other matches D_{i+1}. Evidently, then, D cannot present an impression of continuity of change as it moves through this region, contrary to hypothesis.

To obstruct this reasoning it will not be sufficient to hypothesise that S is infinite; we need specifically that $\langle D_1, \ldots, D_n \rangle$ is densely ordered by temporal succession, that between any pair of stages of D discriminable in respect of $\emptyset$ lies a stage discriminable from them both and from any stage outside the region of D which they flank. We have to suppose that we have in this sense infinite powers of discrimination in D, that we can always directly discern some distinction more minute than any discerned so far. The 'very natural presupposition' earlier referred to is that this is not so.

To summarise our conclusions here: if D is a non-recurrent temporal process of change such that indiscriminability behaves transitively among every selection of stages from it, and if we can directly discern only a finite variety of stages of D – at least in some of its regions – then D must contain seemingly absolutely abrupt changes. Hence if D is everywhere

to give an impression of continuous change, indiscriminability cannot behave transitively among every selection of stages from it; specifically, if D_i and D_{i+1} are adjacent in $\langle D_1, \ldots, D_n \rangle$, derived as above, at least one stage occurring between D_i and D_{i+1} in D must match both.

These considerations are of course incomplete. We have not considered purely spatial processes of change, though one might venture to expect that exactly analogous reasoning would apply; and we have not considered whether there is not a sufficiency condition for seeming-continuity of change in terms of non-transitive indiscriminability. But enough has been done to give this phenomenon a certain dignity. It is not something confined to the psychological laboratory, comparable say, to peoples inability to judge the relative lengths of:

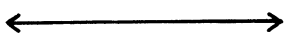

and

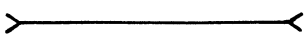

– something of which someone might reasonably require experimental confirmation that he too was subject to it. Granting the presuppositions of the above reasoning, we have rather to regard the non-transitivity of the relation, 'is not discriminably different from', as a reflection of a pervasive structural feature of our sense experience – the continuity of change.

The general lesson then of Example 4 is this. If we attempt to mark off regions of a seemingly-continuous process of change in terms of predicates which are purely observational; – predicates of which it is understood that ostensive definition gives their whole meaning – these expressions are bound to display tolerance in a suitable series of stages selected from the process. In his book on the *Philosophy of Mathematics*, Körner several times characterises perceptual concepts as essentially inexact. This section can be regarded as an attempt to bring this insight into sharp relief. Absolutely any ostensively defined predicate may come to display tolerance, for absolutely anything which it characterises might undergo seemingly continuous change to a point where it could be so characterised no longer. Example 4 is but a tremor, signalling a basic fault, as it were, buried deep in the relation between the nature of our experience and those parts of language by means of which we attempt to give the most direct, non-theoretical expression to it.

This conclusion rests upon two premisses which might be held open to question: that it is right to regard the senses of colour predicates, etc., as purely observational; and that this is a very fundamental fact about their senses, whose sacrifice would be possible only at great cost. For the first no further argument will be provided in this paper. The considerations adduced earlier in the section are surely decisive, provided it is allowed that they are *relevant* – provided the second thesis of the governing view has not been rejected. For the second, however, no argument has so far been presented; we merely voiced concern that 'contact' between language and the empirical world might be attenuated if the use of purely observational predicates was abandoned. No special considerations have been advanced in relation to Example 4 to correspond to the points about convenience, memorability, and social role in relation to the earlier examples. The matter must await the next but one section. First we must consider a general objection to our treatment of all four examples.

IV

To concede that the vagueness of our examples is correctly interpreted as tolerance is to concede that there can be no consistent, non-classical logic for such predicates. But it is natural to suggest that the arguments for this interpretation may have overlooked an essential feature of this sort of predicate: that they typically express distinctions of *degree*. There are degrees of redness, of childishness and, if a smaller heap is regarded as less of a heap, of heaphood also.

When is the distinction between being F and not being F one of degree? Typically, when the comparatives, 'is less/more F than', are in use and when iteration of one of these relations may transform something F into something not F, or *vice versa*. Moreover the semantic relations between the comparatives and the simple descriptions, 'is F' and 'is not F', are such that if a is less/more F than b, then the degree of justice with which a can be described simply as F is correlatively smaller or larger than that with which b can be so described. That is, a twofold classification of possible states of affairs into those which would justify the judgement, 'a is F', and those which would not, misses what is distinctive about a predicate whose application is a matter of degree. For that to be so is exactly for there to be *degrees* of such justice.

It is thus plausible to suppose that a logic for distinctions of this sort cannot be based upon simple bivalence. With such predicates there are, as it were, degrees of truth whose collective structure is that of the set of degrees of *being F*. We shall not speculate what this structure might be, but it does not seem too fantastic to suppose that an indefinitely large or even densely ordered series might be involved, so issuing in infinitely many truth-values.

In this sense it is arguable that our examples do require a special, i.e. non-classical logic. The view that classical logic is inadequate for distinctions of degree is not contested in this paper. What is contested is the idea that the seeming tolerance of the examples is generated by overlooking that the predicates in question express distinctions of degree. This impression views the paradoxical reasoning as essentially depending upon the constraints of bivalence – thus no attention is paid to the point that distinctions of degree are involved. Consider a pair of objects one of which, *a*, we are happy to describe as *F* while *b* is slightly less *F* than *a*. How is *b* to be described? If our admissible descriptions are restricted to '*F*' and 'not-*F*', if we *have* to say one or the other, then presumably we shall describe *b* as *F*. For if something is more like something *F* than something not-*F*, to describe it as *F* is the less misleading of the two alternatives. But the justification with which '*F*' is applied in successive such cases successively decreases. We have no principle of the form: if *a* is *F* and *b* differs sufficiently marginally from *a*, then *b* is *F*; with distinction of degree there are no 'small changes insufficient to affect the justice with which a predicate applies'; they are, on the contrary, small changes *in* the degree of justice with which the predicate may be applied. Of course we do have the principle: if the judgement that *a* is *F* is justified to some large degree and *b* is marginally less *F* than *a*, then the description of *b* as *F* will be better justified than its description as not-*F*. But that is not a paradoxical principle.

Anyone who thinks he at last feels the cool wind of sanity fanning his brow would do well to be clear why we do not still have *this* principle: if *b* is marginally less *F* than *a*, then if the less misleading description of *a* is '*F*', the less misleading description of *b* is '*F*'. Yet if this principle is false there must, in any Sorites type series, be a last case of which we are prepared to say that if we *had* to describe it either as *F* or as not-*F*, the better description would be '*F*'. Why, then, is it usually embarrassing to

be asked to identify such a case without any sense of arbitrariness? Let us say that '*a* is *F*' has a *positive* value just in case '*F*' is a less misleading description of *a* than 'not-*F*'. Then our embarrassment is exactly to identify a last object to which the application of *F* would receive a positive value. But then the suspicion arises that tolerance is with us still; only it is not now the *truth* of the application of *F* that would survive small changes but, so to speak, its positivity.

Is this suspicion justified? One thing is clearly correct about the assumption of bivalence: faced with a situation and a predicate, we have only two choices – to apply or to withhold. There is not a series of distinct linguistic acts in which we can reflect every degree of justification with which a predicate may be applied. The crucial notion to be mastered for practical purposes is thus that of a situation to which the application of *F* is *on balance* justified. Without mastery of this notion, no amount of information about the structure of variations in the degree with which *F* applies entails how the predicate is to be used. Now of this notion may it not still be a feature that it always survives sufficiently small changes? – that if *a* and *b* are dissimilar only to some very small extent, then if describing *a* as *F* is on balance justified, so is thus describing *b*?

It is clear that all the considerations adduced in the previous sections now apply. The introduction of a complex structure of degrees of justification has got us no farther; for among these we have still to distinguish those with which for practical purposes the application of the predicate is to be associated; otherwise we have not in repudiating bivalence done anything to replace the old connection between justified assertion and truth. The distinction in question need not of course be exact, (so one's embarrassment at having to identify a last case to which the application of *F* is preferable to that of its negation is understandable). On the contrary, in the kind of case which we have been considering the concept will be tolerant.

To rehearse the reasons: if we are to be able to *remember* how to apply *F*, then differences too slight to be remembered cannot transform a situation to which its application is on balance justified into one which is not so; if we are to be able to apply *F* just on the basis of *casual observation*, the same applies to differences too subtle to be detected by casual observation; if the distinction between cases to which the application of *F* is on balance justified and others is to be made just on the

basis of how things look, or sound, etc., then any pair of indistinguishable situations must receive the same verdict; finally if F is associated with *moral* or *explanatory* distinctions which we are unwilling to tie to very small changes, we shall likewise be unwilling to allow such changes to generate doubt about the status of a situation previously regarded as on balance justifying description as F. Of course a quite uncritical use is here being made of the notion of a situation to which the application of F is 'on balance' justified. But this is legitimate. As remarked, there must be *some* such notion if a many-valued logic for distinctions of degree is to have any practical linguistic application.

V

We turn to the question whether we could not, at not too heavy a cost, eliminate the tolerance of observational predicates. What is necessary, it seems, is rejection of ostensively defined predicates; hence the initial doubt whether such a purified language could engage with the observational world at all. But, on reflection, it is clear that the dislocation of language and the world of appearance does not have to be as radical as that. When three situations collectively provide a counterexample to the transitivity of indiscriminability, there is nothing occult, as it were, in the circumstance that they do so. It is an observationally detectable difference between indiscriminable situations that one is distinguishable from a third situation from which the other is not; the relation, 'a matches b matches c does not match a', is an observational relation, i.e. one whose application to a trio of objects can be determined just by looking at them, listening to them, etc.

Observational concepts evidently require narrower criteria of reapplication than indistinguishability, if they are to be purified of tolerance. But we should not jump to the conclusion that to provide such criteria will require surrender of observationality altogether, for the phenomenon which is causing the trouble is itself observational. Indeed, the *only* kind of observationally detectable difference which there can be between indiscriminable items is that one should be distinguishable from some third item from which the other is not. So if the class of expressions in question is to remain in contact with observation, we have to look for some form of stipulation which *exploits* the non-transitivity of indis-

tinguishability to provide a basis for describing indiscriminable situations differently. No other explanation can correspond to a distinction which sense-experience can determine to obtain, a distinction which we can simply be shown.

One implication of this suggestion is already apparent. All along our discussion has centred on the use of certain *predicates*. Now there is a tradition, dating at least from Frege, in accordance with which a predicate of individuals is essentially an open sentence of only one individual argument place. The status of a predicate in this sense is indifferent to the occurrence of individual constants and bound variables within the expression, and so to the nature and scope of the procedure required to determine whether the predicate applies. But there is also the narrower traditional conception in whose terms predicates essentially express *properties*, whose application to an individual is a question of scope no broader than the boundaries, so to speak, of the individual concerned; it is this which distinguishes properties from relations. Such was the conception of *quality* to which Locke and Berkeley appealed. Of course it may be appropriate to decide whether a property applies by comparison with cases where it does; but that only serves to emphasise that its application cannot *always* be a matter of comparison.

What is the point of this distinction here? Simply that it is not coherent to demand more than indistinguishability as sufficient for re-application of an observational *quality*. The question whether such a quality is shared by a pair of indistinguishable items cannot essentially turn on what relations they bear to other things; it has to be possible to decide whether an object has the quality by observation of it alone. So the envisaged kind of modification to the senses of observational predicates requires that we abandon their use as expressions of observational qualities. Indeed the interest of the proposal depends upon our abandoning the notion of an observational quality altogether; for if there are such qualities, there can be no objection to introducing predicates to express them.

This is not to say, of course, that language must cease to contain expressions whose syntax is that of a simple predicate, i.e. expressions containing no singular term or quantifier and having but a single individual argument place. But if the conditions of application of such an expression can be determined by observation, they will not be determinable

by observation of a single individual. In a certain sense the world of observation is to be a world of relations. The semantics of any observational predicate will be implicitly relational.

The sense of this last claim becomes clearer if we now take note of a striking aspect of the philosophical psychology of non-transitive matching. Summarily, it does not seem to be possible to conduct experiments with non-transitively matching triads in *memory*. For suppose that a predicate, F, is defined ostensively by reference to some individual, a, which, it is noted at the time, perfectly matches another individual, c; it is understood that F is not to be applied to individuals which match a unless they also match c. Later the trainee comes across b which, so far as he can determine, matches a perfectly; the question is, does b match c? It is evident that the issue is only resoluble by direct comparison, and especially that it cannot be settled by memory, however accurate. For the most perfect memory of c can give no further information than that it looked just like a; which, when non-transitive matching is a possibility, is simply insufficient to determine whether it would match b. This, it should be emphasised, in contrast with our conclusions concerning Example 3, is not a limitation imposed by the feebleness of our memories; it is a limitation of principle.

So, if we are to be able to exercise expressions whose application to matching individuals depends upon their behaviour in relation to a third, possibly differentiating individual, it is clear that we have to be able to ensure the *availability* of the third individual. Expressions of this species will be practicably applicable only in relation to a system of paradigms. Thus we can see, even in advance of attempting a specific stipulation to remove the tolerance of 'red' as displayed in Example 4, that the kind of semantic construction it will have to be is going to tie the application of expressions of colour to the use of a *colour-chart*.

Let us then consider, as a test case, how we might go about the construction of such a chart. What we require of the chart is that it should enable us to identify a last red patch in any series of the type of Example 4. How is this to be achieved? In *The Structure of Appearance*[11] Goodman notes the difficulty that the non-transitivity of matching provides for his concept of a *quale*:

> ... this is somewhat paradoxical; for since qualia are phenomenal individuals, we can hardly say that apparently identical qualia can be objectively distinct.

But Goodman sees a solution:

> ... the fact that some matching qualia are distinct can be accounted for without going beyond appearance; we need only recognise that two qualia are identical if and only if they match all the same qualia.

(The strategy is not original to Goodman. It was also adopted by Russell in *Inquiry into Meaning and Truth* and indeed there described as 'familiar'.)[12]

Let us say that patches *a* and *b* are of the same *Goodman-Shade* just in case they satisfy such a condition; that is, any *c* matches *a* if and only if it matches *b*. Then it would appear that an ideal colour chart – a *Goodman Chart* – would contain samples of all Goodman-Shades, together with instructions about how each is to be described. Evidently, however, we must abandon this ideal straight away. We might chart a region of colour comprehensively in the sense of being confident that any sample of colour which we should intuitively regard as falling in that region would match something on the chart; but we could never have reason to be sure that every Goodman-Shade falling within that region was represented. For there is no way of foreclosing the possibility of finding a sample which is *Goodman-distinguishable* from every sample on the chart from which it is not simply distinguishable; that is, it matches something which such a chart sample does not, or *vice versa.*

It is, indeed, a peculiarity of the notion of a Goodman-Shade that, on a certain natural presupposition, a comprehensive Goodman Chart could not be achieved in principle. Imagine a band of colour varying *uniformly* in the following sense: that there is a constant distance, *d* inches, say, such that we can always distinguish the colours on the band more than *d″* apart but can never distinguish those less than *d″* apart. It follows that the colours at any distinct points on the band are Goodman-distinct, since associated with any pair of distinct points is a point less than *d″* from one and more than *d″* from the other. On such a band the claim of identity of Goodman-Shade at any distinct points is defeasible.

There seems *a priori* no reason why there should not be a band of colour varying uniformly in this sense. But then there can no more be a specific, finite, comprehensive totality of samples of Goodman-Shades than there can be such a totality of points in the linear continuum. The assumption that colour may vary uniformly thus generates a somewhat

startling consequence for the notion of a Goodman-Shade: single Goodman-Shades do not require that regions in which they are exemplified should be extended.

This is, perhaps, paradoxical only if, like Goodman, we suppose that these notions are still in any meaningful sense *phenomenal*. Besides, the kind of comprehensiveness which we were contemplating is not necessary anyway. What is necessary is the construction upon the set of Goodman-Shades of an effective total *order*; such an order should, moreover, coincide with that generated by the intuitive relation of likeness among colour samples in any case where the latter issues a clear verdict. Given a decidable such order, there is no obstruction to stipulating intelligibly a determinate *last* red shade; and then we may in any series like that of Example 4 identify a last sample preceding or coinciding with this shade.

Let us, then, consider what may be achieved in the direction of explaining such an effective order. Plainly, we have to be able to arrange a series of samples in such a way that we can tell of any new sample whether it falls *between* any pair of samples already located; such recognition will of course be crucial in cases where one of the samples is the last red shade. Conversely, if we are given the two outermost flanking samples, it will be sufficient for the task to have an effective notion of betweenness.

When should a shade, *b*, come between *a* and *c* in an intuitively correct arrangement of colours? The notion we really need is that of *b* coming *nearer* to *a* than *c* does; for *b* comes between *a* and *c* just in case *b* is nearer to *a* than *c* is and nearer to *c* than *a* is. Plainly a colour-patch, *b*, should come nearer to *a* than *c* does if and only if it is *more like a* than *c* is. Let us consider then the construction of a Goodman Chart in terms of this basic principle; that sample *b* is to be placed nearer sample *a* than sample *c* is just in case it is more like *a* than *c* is.

The principle has one very clear cut application: among a trio of non-transitively matching samples, the sample which matches the other two must be located between them. And surely there is this much to be said for it: if on a Heaven-sent, correctly constructed Goodman Chart of sufficient degree of refinement of distinction we find a sample *b* located between *a* and *c*, that it is so should be *explicable* by (iterated) application of the principle; if *a* comes to the left of *b*, then there should be something to the left of *a* which *a* matches and *b* does not, and so on.

It is thus plausible to claim that the principle at least gives necessary and sufficient conditions for correctness of mutual location among samples in an existing Goodman Chart. This is not, of course, to say that it can serve as a principle of construction in an effective sense.

We are, moreover, making a very substantial assumption. Suppose that *a* matches *b* matches *c*, while *a* and *c* do not match; what, then, if some *d* matches *a* and *c* but not *b*? The basic principle will require both that *b* and *d* be placed between *a* and *c* and that *a* and *c* be placed between *b* and *d*, contrary to the supposition of total order. It is thus a necessary condition of the possibility of a total order among the relevant Goodman-Shades in terms associating 'is nearer to' with 'is more like to' that this cyclic matching situation should not be a possibility. It will be assumed, since our main conclusions will be unaffected if it does not, that this necessary condition obtains.

Suppose then that we have been given a red and an orange flanking sample, and have advanced, using the basic principle and the kind of ability on which Hume remarked, to a Goodman Chart in which every sample matches just those immediately adjacent to it and no others. Consider a section of the chart:

$$\ldots a;\ b;\ c;\ d;\ e;\ f;\ g;\ h\ldots.$$

Our task is to locate an *x* which matches both *d* and *e*; does it go between them? Evidently in two cases it does: the basic principle puts *x* between *d* and *e* if it matches both or neither of *c* and *f*. But suppose one of the other two cases obtains, for example that *x* matches *c* but not *f*. Still, we may get a decision if *x* matches *b*, for then it must lie between *c* and *d*. But suppose *x* does not match *b*; suppose in fact that *x* and *d* match and are distinguishable from all the same samples located so far. Then all the basic principle tells us is that *x* must lie somewhere in the region between *c* and *e*.

We have no way of locating *x* in relation to *d* on present information; apparently we need more samples on the chart. A sample between *b* and *c*, or between *e* and *f*, matching *x* or *d* but not both would do the trick. Yet we have just noted that the question whether something goes between samples adjacent at this degree of construction is effectively decidable only in two cases of four; so we have no guarantee of being able to locate such additional samples; indeed the spectre is introduced of *reciprocally*

undecidable cases, where the location of x is decidable only by deciding that of y and conversely.

This spectre at least can be banished. The discovery of a y matching x but not d or *vice versa* would actually guarantee a solution to the question of their relative locations in the region between c and e.[13] Still we cannot convincingly claim to have an effective principle of construction, since we have no effective means of finding such a y.

There is a yet more radical deficiency. If x and d are as a matter of fact of the same Goodman-Shade, then not only have we no effective way of determining that they are; we have no way of recognising the fact *at all.* This is a fundamental point. Goodman, we noted earlier, congratulated himself on being able to explain a distinction between matching and identity among qualia without going beyond the realm of appearance. Certainly, the notion of Goodman-identity is explained by means of a phenomenal relation and a quantification over ostensibly phenomenal individuals. Yet, evidently, it is only in a formal sense that we have not gone beyond appearance in this account; for, as Russell noticed, the resulting relation is not one which may *appear to obtain.*

The concept of a Goodman-Shade is not merely not an observational notion; it is a transcendental notion. Nothing counts as discovery of identity of Goodman-Shade. This casts serious doubt on the suitability of the notion to serve in the construction of a semantic rule. For, surely, sense can be given to an expression only by reference to conditions whose satisfaction we can determine at least in principle. After a stipulation of the kind considered, the ability to recognise whether something was red would sometimes require the ability to recognise identity of Goodman-Shade, viz. when what we have is in fact a sample of the last red shade. But this is an ability which we do not have.

There is no effectively decidable order among Goodman-Shades based on likeness. *A fortiori* there is no effective stipulation of the sort considered: a stipulation based on the notion of a Goodman-Shade and providing a sharp, decidable red/orange distinction in any series of the type of Example 4. The concept of the last red shade is underdetermined, for nothing amounts to recognition of which shade it is. Indeed Goodman's whole strategy for surmounting the difficulty of nontransitive matching amounts to nothing other than the introduction of a spuriously phenomenal identity to which nothing in our experience can correspond.

Not that Goodman introduced the notion with a view to providing the kind of stipulation which we are looking for; but we must look elsewhere.

Could we not instead simply devise *ad hoc* paradigms? Consider a colour-chart complete in the sense distinguished earlier; i.e., we are confident that anything we should wish to regard as falling in the red/ /orange region will match something on the chart. It is plausible to suppose that likeness provides an effective order among any finite set of colour samples, so suppose the chart samples to be so arranged. Then we can generate a decidable red/orange distinction as follows. Select some patch towards the middle; then any colour patch matching something on the chart either matches the selected patch or it does not; if it does, it is red; if it does not but matches a sample to the left of the selected patch, it is red; otherwise it is orange.

Naturally we could not guarantee that duplicates of a chart would always deliver the same verdict. Charts could look absolutely similar, and even satisfy the condition that the *n*th sample on either matched and was distinguishable from exactly the same samples on the other chart as its own *n*th sample, yet deliver discrepant results. But they would not often do so. Besides, the situation is nothing new. Rulers, for example, sometimes give different results. A final criterion for one system is deposited in Paris; and we could do the same with a colour-chart.

Nevertheless the generalisation of this proposal seems quite ludicrous in practical terms. We are confronted with the imaginary spectacle of a people quite lost without their individual wheelbarrow loads of charts, tape-recordings, smell- and taste-samples and assorted sample surfaces. But this does the proposal an injustice. There will be no need for all this portable semantic hardwear. To pursue the analogy with the use of rulers: it is true that, without a ruler, we can only guess at lengths; but after the introduction of an *ad hoc* paradigm for colours, the use of colour predicates will presumably be analogous not to that of expressions like, 'two feet long', but rather to that of expressions like, 'less than two feet long', i.e. expressions of a *range* of lengths. Of such expressions the criterion of application is still measurement; but unless the case is peripheral, we can tell *without* measuring what the outcome of measurement would be. Training in the use of paradigms might be essential if one is to grasp the sense of such expressions; but, once grasped, most cases

of practical application could be decided without the use of paradigms; for most practical purposes the wheelbarrow could be left behind.

It appears, then, that were we to adopt such stipulations as a general strategy, it would not have to affect our use of observational language very much at all. At present we can tell of anything red that it is so just by looking at it. This would still usually be true after the proposed stipulation; and if the new distinction was suitably located, cases where it was not true could generally coincide with borderline-cases of the old red/orange distinction. The use of predicates so refined could thus greatly resemble their present use; the distinctions which they expressed would be empirically decidable; and there would be one crucial disanalogy – they would be tolerance-free.

It is, indeed, apparent that exactly parallel considerations may be brought to bear upon our earlier treatment of Examples 1 and 3. Even after a precise re-definition of 'heap', we would be able to learn to tell in most cases just by casual observation what verdict the new criterion would give if applied; it would seldom be necessary actually to count the grains. And the distinction between red and orange, supposing an exact distinction were drawn by means of a chart, would be unmemorable only within that small range of shades which could not by unaided memory be distinguished from the last red sample. It would thus appear that the cost of eliminating tolerance in cases of these two types need not after all be high, since we could expect to be able to tell in general just by looking at, etc., an item on which side of the dividing line it would fall.

If there need after all be no substantial sacrifice in endowing formerly observational predicates with exact boundaries, what has become of the alleged profound tension between phenomenal continuity and language designed to express how things seem to us? The answer is that we have simply swept it under the carpet. The possibility of our dispensing with paradigms for most practical purposes depends upon our capacity e.g. to distinguish between cases where we could tell whether or not 'red' applied just by looking and cases where we could not, where we should have recourse to a chart. If we are able to make such a distinction, what objection can there be to introducing a predicate to express it? But then, it seems, the semantics of this predicate will have to be observational. For on what other basis should we decide whether something looks as though comparison with a chart would determine it to be red than how

it looks? Of any pair of colour patches which look exactly alike, if either looks as though the chart would deliver the verdict, 'red', both must. So the new predicate, introduced to reflect our capacity to make this distinction, will be applicable to both members of any pair of matching colour-samples if to either.

Of course, there is no reason to have any such predicate; but equally there is no reason not to. If we were sometimes able to tell without using a chart whether something is red, it would surely be possible to make intelligible to us a predicate designed to apply in just such circumstances. A language all of whose observational concepts were based on paradigms would avoid containing tolerant predicates only by not containing means of expression of all the observational distinctions which we are in fact able to make. Naturally it would make no difference to this point if we insisted on actually using the paradigms in every case, and hang the inconvenience, as it were. We *could* dispense with them almost always, so we should just be insisting on a charade.

Would it not, moreover, be quite absurd to propose that the tolerance of such new predicates – 'looks as though it would lie to the left of the last red shade', 'looks as though it contains fewer than ten thousand grains', etc. – might in turn be stipulated away? Their meaning will not permit it; it cannot be allowed of things which look exactly alike that one may look as though it satisfies some condition which the other looks as though it does not – unless how a thing looks may not be determined by looking! But the earlier treatments of Examples 1 and 3 involved overestimation of our interest in preserving the tolerance of the predicates concerned only if we posses a coherent understanding of these new predicates; if Examples 1 and 3 do not, after all, pose a substantial problem for the governing view, it is because of our capacity to handle expressions falling within the scope of Example 4.

VI. CONCLUSION

Let us, then, review the character of the difficulty which Example 4 would appear to pose for the governing view.

It is a fundamental fact about us that we are able to learn to classify items according to their appearance, that we are able consistently, as it seems to us, to apply or to withhold descriptions just on the basis of how

things strike the senses. In a discrete phenomenal world there would be no special difficulty – no difficulty, that is, not inherent in the idea of a semantic rule as such – in viewing our use of such expressions as essentially nothing but the following of rules of which it was a consequence that indiscriminable phenomena should receive the same description. But if mutually exclusive use is made of a pair of such predicates, and if cases to which one applies permit of continuous transformation into cases where the other applies, it cannot be correct to represent the use made of either predicate just as the doing of what is required by a set of rules with such a consequence. Yet we are constrained – if the relevance is allowed of considerations to do with what we should regard as adequate explanation of such expressions, or with certain criteria which we should accept of misunderstanding such an expression – to attribute to the rules governing these predicates precisely such an implication; and all the phenomena which we confront in our world impress us as capable of continuous variation.

In the Introduction, the difficulty was presented, starkly, as that of the inadequacy of any inconsistent set of rules to explain a consistent pattern of behaviour. This needs a little refinement. It is, to begin with, unclear how far our use of e.g. the vocabulary of colours *is* consistent. The descriptions given of awkward cases may vary from occasion to occasion. Besides that, the notion of using a predicate consistently would appear to require some objective criteria for variation in relevant respects among items to be described in terms of it; but what is distinctive about observational predicates is exactly the lack of such criteria. So it would be unwise to lean too heavily, as though it were a matter of hard fact, upon the consistency of our employment of colour predicates. What, however, may be depended upon is that our use of these predicates is largely *successful*; the expectations which we form on the basis of others' ascriptions of colour are not usually disappointed. Agreement is generally possible about how colours are to be described; and this, of course, is equivalent to saying that others *seem* to use colour predicates in a largely consistent way.

It is of this fact which the governing view can provide no account. A semantic rule is supposed to contribute towards determining what is an acceptable use of its associated expression. The picture invoked by the first thesis of the governing view is that there is, for any particular

expression in the language, a set of such rules *completely* determinant of when the expression is used correctly; such a set thus provides a model of the information of which a master of the use of the expression may be deemed to be in possession. Clearly, however, the feasibility of such a picture requires that the rules associated with an expression, about whose use we generally agree, be consistent. For if they issue conflicting verdicts upon the correctness of a particular application of the expression, it cannot be explained just by appeal to the rules why we agree that the application is e.g. correct.

The problem presented for the first thesis by the occurrence of tolerant predicates, or of any other kind of semantically incoherent expression, is not that, in a clear-cut way, nothing can be done to implement an inconsistent set of instructions. It is true that, strictly, anything that is done will conflict with a part of them. But we can imagine a game whose rules conflict but which is nevertheless regularly and enjoyably played to a conclusion by members of some community because, for perhaps quite fortuitous reasons, whenever an occasion arises to appeal to the rules, the players concur about which element in the rules is to be appealed to, so that an *impasse* never comes about. We need not enquire whether they have noticed the inconsistency in the rules. The point of the analogy is that in practice they always agree whether a move is admissible, as we generally agree whether something is red. The analogue of the first thesis in relation to this example is the notion that the rules completely determine when a particular move is admissible. But while it may be true that the authority of the rules can be cited for any of the moves the community actually makes, it is plain that the rules alone do not provide a satisfactory account of the practice of the game. For someone could master the rules yet still not be able to join in the game, because he was unable to guess what sort of eclectic application of them an opponent was likely to make in relation to any given move.

An outsider attempting to grasp our use of a tolerant predicate would presumably not encounter exactly this difficulty; it would be clear that we were not prepared to allow remote consequences of its tolerance, inferred by means of reasoning of the Sorites type. But the difficulty of principle for the first thesis is the same. The rules of the game do not provide an account of how the game is played, for it is possible that someone might grasp them yet be unable to participate. The semantic

rules for an expression are supposed to provide an account of its correct use; they cannot do so if someone whose use of it differed radically from ours could still be thought of as in possession of exactly the same brief – as he can be if it consists in an inconsistent set of instructions.

The comparison of language with a game is in many ways an extremely natural one. What better explanation could there be of our ability to agree in our use of language than if, as in a game, we are playing by the same rules? So we are attracted towards the assimilation of our situation to that of people to whom the practice of a game has been handed down via many generations but of which the theory has been lost. Our task is to infer the theory. The burden of this paper, however, has been that this image is in conflict with the principles of investigation into the semantics of an expression which we find it natural to allow. These principles yield an account of the senses of certain vague predicates which pushes our agreement in their application beyond explanation by appeal to what the rules for their use require. And it is doubtful whether an intuitively satisfactory conception can be achieved of what an investigation into the semantics of an expression might be which did not admit these priciples, – unless the notion is abandoned that such an investigation is something which only a master of the investigated language is optimally placed to carry out.

Such a conclusion would seem to force on us a more purely behaviouristic concept of how a theory of language-use should be accomplished, and a corresponding shift in the concept of a semantic rule. There would no longer be any room for the idea that such rules might be discoverable by means of such a sort that our use of the associated expression(s) could prove to be in conflict with or otherwise inexplicable by reference to them; there would no longer be any room for the idea that by reflection on the kind of training which he has received in the use of some expression, the criteria which he would employ to judge that someone misunderstands it, his concept of the purpose or interest of the classification which it effects, or his awareness of his own intellectual and perceptual limitations, a speaker of a language has access to sound conclusions about his understanding of an expression which a mere observer of his use of it has not.

This would be one response to the difficulties generated by the richer, more natural methodology of the governing view. It can be expected to

encounter severe difficulties of its own: in particular, it is unclear that an adequate characterisation can be achieved from so restricted a standpoint of in what the vagueness of the sorts of example with which we have been concerned *consists*. But if the notion of a semantic rule is not to be abandoned altogether, *some* more restricted account of the epistemology of semantic rules is required than that afforded by the governing view.

All Souls College, Oxford

NOTES

* This discussion was largely motivated by two unpublished papers: Aidan Sudbury's 'Are Imprecise Terms Essential?' and Michael Dummett's 'Wang's Paradox'. The paper has benefitted by extensive discussion with both Sudbury and Dummett.

[1] Cf. *Grundgesetze*, Vol. II, Section 65.

[2] Elsewhere Russell takes the vagueness of ordinary language more seriously; e.g. in 'Vagueness', *Austral. Jour. of Psych. and Philosophy* **1** (1923), 84–92, the notion that vagueness is a flaw is tempered by pessimism about its remediability.

[3] In the *Structure of Appearance*.

[4] Traditionally, attributed along with its variant, the 'Bald Man' to Eubulides. (See e.g. Diogenes Laertius, Lives, ii, 108.)

[5] 'On the Ultra-intuitionistic Foundations of Mathematics', in *Infinitistic Methods*, Pergamon Press, Oxford, 1961, pp. 201–223.

[6] Esenin-Volpin's interest in the example derives from the possibility, which it suggests, of satisfying orthodox arithmetical postulates for the successor function in an open-ended but finite domain. The existence of such domains depends upon the admissibility of the concepts of which they are the apparent extensions; that is, it depends upon the correctness of the view that we may have a coherent understanding of predicates whose semantics *are* strictly incoherent. For unless a predicate of the genre in question really does have an incoherent semantics, it will not yield a structure in which the relevant postulates are satisfied; but unless such predicates are acceptable, we shall not be much interested in the pure mathematics of their extensions. There will be, unfortunately, no further discussion in this paper of the significance of these examples for the concept of infinity.

[7] We spare the reader the irrelevant complications involved in how such a stipulation would apply to substances of varying granular coarseness, viscosity, etc.

[8] Intuitively, a colour predicate is basic in a given language if it does not express a shade of a more general colour for which means of expression exists in the same language. Obviously, as it stands, the notion is fragile.

[9] As they would sometimes be, if the distinctions were made in terms of precise numbers of heartbeats.

[10] There is no presupposition here that a definitely correct verdict can always be reached; but if it cannot, that in turn must be the situation with respect to each item in question.

[11] Chapter IX, Section 2.

[12] Chapter VI.

[13] *Proof:* Let y match d but not x, and e but not c. Then, letting 'M' = 'matches', '$\bar{M}$' = 'does not match', '$a\ b\ c$' = 'b is between a and c':

(i) $y\ M\ d, d\ M\ c, y\ \bar{M}\ c$; $\therefore\ y\ d\ c$; hence y is right of d.

(ii) $y\ M\ d, d\ M\ x, y\ \bar{M}\ x$; $\therefore\ y\ d\ x$; hence by (i), x is left of d.

The case where y matches c but not e is quite similar.

Suppose on the other hand that y differentiates x and d the other way about, i.e., that y matches x but not d.

Let y again match e but not c. (The converse case is again essentially similar.) Then:

(i) $y\ M\ e, e\ M\ d, y\ \bar{M}\ d$; $\therefore\ y\ e\ d$; hence y is right of d.

(ii) $y\ M\ x, x\ M\ d, y\ \bar{M}\ d$; $\therefore\ y\ x\ d$; hence by (i), x is right of d.

So if a differentiating y matches only one of c and e, the question of the location of x in relation to d is soluble.

What, though, if y matches both or neither of c and e? Somewhat surprisingly the basic principle, conceived as determining a total order, proves strong enough to rule out both cases. Again, let y match d but not x:

Suppose $y\ M\ c, y\ M\ e$; then:

(i) $y\ M\ e, e\ M\ x, y\ \bar{M}\ x$; $\therefore\ y\ e\ x$

(ii) $y\ M\ c, c\ M\ x, y\ \bar{M}\ x$; $\therefore\ y\ c\ x$.

Hence both c and e lie between y and x; so x cannot lie between c and e, contrary to the hypothesis that, save for y, x matches all and only the patches matched by d, (whence $c\ x\ e$).

Suppose alternatively $y\ \bar{M}\ c, y\ \bar{M}\ e$; then:

(i) $y\ M\ d, d\ M\ e, y\ \bar{M}\ e$; $\therefore\ y\ d\ e$.

(ii) $y\ M\ d, d\ M\ c, y\ \bar{M}\ c$; $\therefore\ y\ d\ c$.

Hence y is required, absurdly, to lie both right and left of d.

Now let y match x but not d:

Suppose $y\ M\ c, y\ M\ e$; then:

(i) $y\ M\ c, c\ M\ d, y\ \bar{M}\ d$; $\therefore\ y\ c\ d$.

(ii) $y\ M\ e, e\ M\ d, y\ \bar{M}\ d$; $\therefore\ y\ e\ d$.

Hence, absurdly, y is required to lie both left of c and right of e.

Suppose $y\ \bar{M}\ c, y\ \bar{M}\ e$; then

(i) $y\ M\ x, x\ M\ c, y\ \bar{M}\ c$; $\therefore\ y\ x\ c$.

(ii) $y\ M\ x, x\ M\ e, y\ \bar{M}\ e$; $\therefore\ y\ x\ e$.

Here (i) and (ii) are conjointly satisfiable only if both y and x lie to the left of the region between c and e; but x must lie *within* that region since, prior to the discovery of y, it matched and was distinguishable from the same samples as d.