account, though not with Miss Anscombe's, that I may sometimes have a practical problem about *whether* to do what I want to do. This would be true of those cases in which I feel that what I want to do is wrong. In such cases I need not cease to want it (indeed, this may be psychologically impossible for me), but only to realize that I'm not morally free to act on this want. Case A is an example of this. (And in case B, could the felons take advantage of Miss Anscombe's doctrine to plead in court that after all they had permitted me to do what I wanted?)

The argument of (i)–(vi) above goes wrong, then, in step (v). If an agent, in S, decides that a is the preferable alternative all things considered, and solves his problem by setting himself to do a, then what follows is that in S he is most willing to do a — not that he wants to do it. If this seems like a small inference from the preceding steps, it is. Steps (i)–(iv) are superfluous, and step (vi) is false for the same reason that (v) is false. Even granting that (i)–(iv) are all true (which is not above suspicion), (v) and (vi) do not follow.

The positive result of this is not very great; at most, it is a suggestion that one's wants play more varied roles in conduct determination than any simple thesis can account for, and that practical problems are solved in any number of ways. But the negative result is important, too: that someone's voluntarily setting himself to do something is not necessarily a sign that he wants to do it.

Received May 15, 1967

Policing the Aufbau

by DAVID K. LEWIS

UNIVERSITY OF CALIFORNIA AT LOS ANGELES

Introduction

CARNAP's *Aufbau*[1] sketches a remarkably ambitious construction. Given just one primitive phenomenal relation, he seeks to define enough concepts to provide a language adequate for all of science.

The constructed concepts are supposed to be coextensive with certain familiar ones. The *Aufbau* is commonly dismissed as a failure because discrepancies would appear under unfavorable circumstances. That verdict is

premature. If there are few discrepancies under actual circumstances, the constructed concepts might be just as adequate for science as the familiar ones they approximate and replace. A mere chance of discrepancies is too bad, but not fatal. It would take frequent discrepancies to spoil the construction, by Carnap's own standards of success. The frequency has not been much investigated — understandably so; the needed computing power has become available only lately. Therefore it remains an open question whether the *Aufbau* succeeds or fails on its own terms.

If the original construction gives too many discrepancies, perhaps a more elaborate version would work better. The version presented, after all, is tentative; indeed it serves mostly as an illustration for Carnap's general discussion of logical constructions. There is plenty of room for improvement.

In particular, there are unexploited opportunities to police the construction. Spurious instances of a constructed concept, instances which do not fall under the familiar concept being approximated, often turn out to behave strangely later in the construction. Thereby they may be recognized as suspect and removed. I shall show how this tactic might be used in the early stages of the construction to help fight the difficulty Goodman calls "imperfect community."

Quality Classes and Their Similarity

Carnap begins his construction with a primitive relation of part similarity (and temporal precedence) between *elementary experiences*, momentary slices of one's total stream of experience (sections 108–10). To speak outside the system, as we must to explain its primitive, two elementary experiences are *part similar* just in case some constituent in one resembles some constituent in the other. There might be a crimson spot at the center of the visual field in one and a scarlet spot slightly off-center in the other, or a smell of skunk in one and a smell of burning rubber in the other, or fright in one and an edgy feeling in the other.

A *similarity circle* (of part similarity) is any maximal class of elementary experiences connected pairwise by part similarity, so that each is part similar to each (section 111). To speak outside the system, a similarity circle is supposed to contain just those elementary experiences in which there occur one or more members of some maximal class of mutually similar constituents. One similarity circle might contain just those elementary experiences in which there occur color spots within a certain small area of the visual field and within certain small ranges of hue, brilliance, and saturation; or just those in which there occur tones within certain small ranges of pitch, loudness, and timbre; or just those in which there occur fright, edgy feelings, or other like emotions.

Unfortunately, not all possible similarity circles are of this intended sort. Spurious similarity circles may appear, connected by a fortuitously complete set of part similarities in miscellaneous respects. Similar centrally located greenish spots may occur in elementary experiences E_1 and E_2, similar lower left brownish spots in E_2 and E_3, similar skunk smells in E_3 and E_4, similar cello tones in E_4 and E_1, similar frights in E_1 and E_3, and similar flute tones in E_2 and E_4. This coincidence ought not to unite E_1, E_2, E_3, and E_4 in a similarity circle, yet it does. Here is the difficulty of imperfect community.

A *quality class* is any maximal subclass of a similarity circle which is wholly included in every similarity circle containing at least half its members (section 112).[2] To speak outside the system, a quality class is supposed to contain just those elementary experiences in which some one constituent occurs. One quality class might contain just those in which crimson of a certain hue, brilliance, and saturation appears at a certain place in the visual field. Hence a similarity circle is supposed to be a union of the quality classes which correspond to several mutually similar constituents. Quality classes are the classes carved out by intersecting similarity circles (ignoring the relatively small intersections due to overlap of the quality classes themselves), and they have been defined accordingly.

Within the system, quality classes could not have been defined in terms of the constituents occurring in elementary experiences; no such things are recognized to exist. Rather the quality classes themselves take the place of constituents; for this reason they are called *quasi constituents* of the elementary experiences, and their construction is called a *quasi analysis*.

Spurious similarity circles will usually yield spurious quality classes: classes of elementary experiences which, as we say outside the system, share no common constituent. We shall see how to distinguish these spurious quality classes from genuine ones by means available within the system itself.

Two quality classes are *similar* just in case each elementary experience in one is part similar to every elementary experience in the other; that is to say, just in case both are included in some similarity circle (section 114). Speaking outside the system, quality classes are supposed to be similar just in case they correspond to similar constituents. But not all possible similarity is of this intended sort. Two quality classes may be spuriously similar by both being included in some spurious similarity circle.

A *sense class* is any maximal class of quality classes which is connected throughout by chains of similarities (section 115). There is supposed to be one sense class for each sense modality: for instance, a visual sense class containing just those quality classes which correspond to color spots. These sense classes are multidimensional arrays. The visual sense class, for in-

stance, has five dimensions: two visual field coordinates, hue, brilliance, and saturation.

A Method of Identifying Spurious Quality Classes

Any genuine quality class will take its place in a sense class. There it will have many neighbors: genuine quality classes, genuinely similar to it. Some, at the edges of their sense classes or near gaps where their neighbors have failed to be constructed, will have fewer neighbors than the rest. Still, almost all genuine quality classes will be genuinely similar to many other quality classes. Genuine quality classes may also be spuriously similar to other quality classes, spurious or genuine.

A spurious quality class, on the other hand, cannot be included in any but spurious similarity circles. So it can be spuriously, but never genuinely, similar to other quality classes, spurious or genuine. But there is no evident reason why it should be more susceptible to spurious similarity than a genuine quality class. I conclude that it is likely to be similar to relatively few other quality classes, or to none.

Let the *neighborliness* of a quality class be the number of quality classes to which it is similar. Let a *hermit* be any quality class with low neighborliness; more precisely, any quality class such that no more than a specified small fraction of all other quality classes are less neighborly than it is. This fraction is the fraction of hermits among all the quality classes. If it were set at .1, say, the hermits would be defined as the least neighborly tenth of the quality classes.

I have argued that if the fraction is set properly, then we can expect almost all and only hermits to be spurious. So if we purge the hermits from the population of quality classes we will get rid of few genuine quality classes and most spurious ones. That is my proposal.

Variants of the Method

We might want to guard strongly against purging genuine quality classes. We might set a low fraction of hermits, or we might even decide to remove only *total hermits*: quality classes with neighborliness 1, similar only to themselves. This is the most cautious version of my proposal. If spurious similarity is rare, plenty of spurious quality classes might be total hermits. But few genuine quality classes — even corner members of low-dimensional sense classes — would ever lose all their neighbors.

Next, notice that the fewer spurious similarities there are present, the more sharply do genuine and spurious quality classes contrast in neighborliness, so the more improvement we can get by removing hermits. It would be helpful to identify and remove spurious similarities before we identify

and remove hermits. We can do so as follows. Let the *bypass distance* of a similarity between quality classes A and B be the number of steps in the shortest chain of similarities from A to B which would remain after the similarity between A and B had been removed. (If no chain would remain — that is, if the similarity connects what otherwise would have been two sense classes — let the bypass distance be infinite.) Let a *shortcut* be any similarity with high bypass distance; more precisely, any similarity such that no more than a specified small fraction of all other similarities have higher bypass distance than it has. (Let a *total shortcut* be any similarity with infinite bypass distance.) Since spurious similarities occur at random, they will often connect quality classes which would otherwise be far apart, perhaps even in different sense classes. So if the fraction of shortcuts is set properly, then we can expect almost all and only shortcuts to be spurious.

Therefore we might police the construction in two stages: first remove those similarities which are shortcuts (or total shortcuts, if we want to be cautious), then remove those quality classes which are hermits. Or we might go deeper. Instead of removing the shortcuts and hermits themselves, we might identify and remove those similarity circles which turned out to yield shortcuts or hermits, and afterwards use the surviving similarity circles to repeat the construction of quality classes and similarities.

Finally, I have been supposing so far that the fractions of shortcuts and hermits have somehow been chosen outside the construction. But we might let them be determined within the construction, making them depend on the statistics of bypass distance and neighborliness. We might take the highest bypass distance for which frequency in the frequency distribution of bypass distances reaches a local maximum, and let the shortcuts be those similarities with bypass distance falling more than a certain amount above that. Likewise we might take the lowest neighborliness for which frequency in the frequency distribution of neighborliness reaches a local maximum, and let the hermits be those quality classes with neighborliness falling more than a certain amount below that.

Received May 1, 1967

NOTES

[1] Rudolf Carnap, *Der Logische Aufbau der Welt* (Berlin and Schlachtensee: Weltkreis-Verlag, 1928). English translation by Rolf A. George, *The Logical Structure of the World* (Berkeley and Los Angeles: University of California Press, 1967). See also Nelson Goodman's exposition and criticism of the *Aufbau* in *The Structure of Appearance* (Cambridge, Mass.: Harvard University Press, 1951), Chapter V.

[2] This is not quite Carnap's definition. I have made the correction proposed by Goodman in *Structure*, Chapter V, section 5. I have also required every quality class to be included in some similarity circle, since otherwise certain unions of intersections of overlapping quality classes would satisfy the definition.