VAGUENESS AND INDUCTIVE MOLDING

John R. Welch

This is a revised version of "Vagueness and Inductive Molding," Synthese 154 (2007), 147–172.

ABSTRACT

Vagueness is epistemic, according to some. Vagueness is ontological, according to others. This paper deploys what I take to be a compromise position. Predicates are coined in specific contexts for specific purposes, but these limited practices do not automatically fix the extensions of predicates over the domain of all objects. The linguistic community using the predicate has rarely considered, much less decided, all questions that might arise about the predicate's extension. To this extent, the ontological view is correct. But a predicate that applies in some contexts can be reasonably extended to other contexts where it is initially vague. This process of development approximates the cognitive remedy for vagueness that the epistemic view prescribes. The process is piecemeal and inductive, akin to what von Wright described as the molding of concepts.

Vagueness cannot be understood apart from the backdrop of classification, for vagueness is classification gone awry. Hence these pages explore the classification of particulars, both its clear successes and vague failures. How we classify unique particulars is the theme of Sections 2 and 3, which are primarily descriptive. Section 2 identifies a way of classifying particulars that pervades discourse of all sorts, and Section 3 illustrates its use in a field notorious for vagueness: ethics. Why a certain particular should (or should not) be classified in a certain way is a normative question, however, and it occupies Sections 4 and 5. Section 4 proposes a norm for cogent arguments by analogy, and Section 5 illustrates how the norm might resolve vagueness in one kind of ethical dispute. This norm, which has a strong probabilistic component, is one way of affirming that probability is a guide to life.

1. Introduction

Vagueness is epistemic, according to some (Sorensen 1988; Williamson 1994). Objects have clear boundaries, and vague language reveals ignorance of these boundaries. Vagueness is ontological, according to others (Tye 1990). Objects like Mount Everest are inherently fuzzy, and vague language reflects the underlying fuzz.

This paper deploys what I take to be a compromise position. It discusses the central case of vague predicates, though adjectives, adverbs, quantifiers, definite descriptions, and proper names may also exhibit vagueness. Predicates are coined in specific contexts for specific purposes, but these limited practices do not automatically fix the extensions of predicates over the domain of all objects. The linguistic community using the predicate has rarely considered, much less decided, all questions that might arise about the predicate's extension. To this extent, then, I take the ontological view to be correct: there may be no fact of the matter of whether a man with 229 hairs on his head is bald. But this is not the end of the matter. A predicate that clearly applies in some contexts can be reasonably extended to other contexts where it is initially vague. This process of development approximates the cognitive remedy for vagueness that the epistemic view prescribes.

In addition, this developmental process appears to be what von Wright was groping for in his reflections on molding concepts (von Wright 1963, vii, 5, 138, 171). The urge to undertake conceptual investigation is one of the main reasons for doing philosophy, according to von Wright. This urge, he says, arises from bewilderment about the meaning of words. This is not the type of bewilderment produced by unfamiliar terms. It arises in connection with familiar terms when we do not know the grounds appropriate for their use. The aim of this type of conceptual investigation

> … is not to "uncover" the existing meaning (or aspect of meaning) of some word or expression, veiled as it were behind the bewildering complexities of common usage. The idea of the philosopher as a searcher of meanings should not be coupled with an idea or postulate that the searched entities actually are there—awaiting the vision of the philosopher. If this picture of the philosopher's pursuit were accurate, then a conceptual investigation would, for all I can see, be an empirical inquiry into the actual use of language or the meaning of expressions.
>
> Philosophical reflexion on the grounds for calling a thing 'x' is challenged in situations, when the grounds have not been fixed, when there is no settled opinion as to what the grounds are. The concept still remains to be moulded and therewith its logical connexions with other concepts to be established. The words and expressions, the use of which bewilder the philosopher, are so to speak in search of a meaning. (von Wright 1963, 5)

In the spirit of these remarks, what I propose in this paper is a strategy for conceptual investigation. The basic idea is to mold concepts and thereby reduce vagueness through a process of inductive inference. 'Inductive', as I will be using the term, will mean 'nondemonstrative' unless otherwise noted. The argument is modest in that I will not claim that inductive molding can eliminate vagueness. The law of excluded middle does not always hold. However, I will argue that the truth-value gaps associated with these failures need not be permanent, that they can be reduced on a piecemeal basis. The engine of reduction, I claim, is inductive logic.

Vagueness cannot be understood apart from the backdrop of classification, for vagueness is classification gone awry. Hence these pages explore the classification of particulars, both its clear successes and vague failures. How we classify unique particulars is the theme of Sections 2

and 3, which are primarily descriptive. Section 2 identifies a way of classifying particulars that pervades discourse of all sorts, and Section 3 illustrates its use in a field notorious for vagueness: ethics. Why a certain particular should (or should not) be classified in a certain way is a normative question, however, and it occupies Sections 4 and 5. Section 4 proposes a norm for cogent arguments by analogy, and Section 5 illustrates how the norm might resolve vagueness in one kind of ethical dispute. This norm, which has a strong probabilistic component, is one way of affirming that probability is a guide to life.

2. Core Classification

About 150 Yanomama Indians eke out a Stone Age existence deep in the Amazon rain forest, in an isolated village on the border between present-day Venezuela and Brazil. A team of anthropologists and journalists visited the area some years ago, and one of the journalists reported the following encounter:

> The little bull of a man with brushcut hair and only a bark string around his waist was studying our Venezuelan Air Force Super Puma helicopter like a scientist. Once before he had seen something similar, and he drew his arm high across the sky in an arc. "I kept calling him to come down, but no luck." What about the chopper, we asked, what was it? He paused for a moment, then tentatively offered: "It's an animal, a hashimo"—a smooth-feathered green grouse—and by way of explanation waved his arms and made the thrashing sound of a big bird exploding from the underbrush. What kept it here? "It's a pet." Did he want to go for a ride? "Maybe later. A long time from now. A really long time." (Reiss 1990, 46)

Of the several remarkable features of this encounter, I would like to focus on just one: Yokokoma's classification of the helicopter as a hashimo. How did he manage to do that? And how do we manage to classify it otherwise?

Perhaps Yokokoma thought that since the smooth, green, flying thing he saw that morning was a hashimo, and since this is smooth, green, and flying, this is therefore a hashimo. Simplifying and schematizing, we might understand his reasoning thus:

C_1: That F is G.
This is F.
Hence this is G.

Here the first premise, 'That F is G', is what I will call a simple categorical. Where do such premises come from? Clearly they are compounds, conjuncts of classifications like 'That is F' and 'That is G'. These bare, compact identifications are then all there is to the premises of C_1, for the first premise conjoins them, and the second premise is one. Suppose we call these rudimentary orderings of the world with the form 'δ is E', where 'δ' stands for a demonstrative pronoun or a proper name and 'E' for a class term, core classifications.

Yokokoma may have been up to something more complicated, however. He may have reasoned that since every smooth, green, flying thing is a hashimo and this is smooth, green, and flying, this is therefore a hashimo. Simplifying as before, we might credit him with an inference of the following form:

C_2: All F are G.
This is F.
Hence this is G.

How might classifiers like Yokokoma and ourselves arrive at the universal first premise? Perhaps we know that all F are G simply because we know the language; the premise may be a

linguistic truth. Alternatively, it may be an inductive generalization. We may have examined some or all of the Fs, verifying in each case that this F is G, and thus are prepared to generalize about all Fs being G. Premise one would then be a conclusion inferred entirely from premises like 'This F is G'. In that case, the component sentences of C_2 would all be either isomorphic to the core classification of premise two or justified by such sentences.

C_1 and C_2 are not atypical. Successfully getting around in the world requires us to identify food, danger, and potential mates, and these identifications require practice and proficiency in core classification. How we core classify is therefore a question of cardinal importance. Though the question can be addressed from many points of view, the biggest part of the answer can be put in the fewest words. Positive core classification is by analogy. The astronomer identifies a quasar, the camper a lichen, the musician a half-tone by perceiving the similarities between a unique object and previously cognized quasars, lichens, and half-tones. In all these cases, there is the known old, the unknown new, and the assimilation of the latter to the former through the relation of similarity. This assimilation is all but transparent in ordinary English expressions such as 'This looks like a mantis' and 'That smells like fire'. Negative core classification, on the other hand, is by disanalogy. The numismatist's judgment that this is not gold arises from the known old, the unknown new, and the perceived dissimilarity between the two. Hence I propose the following thesis: positive core classification is by analogy; negative core classification is by disanalogy. Call this the analogy thesis for short.

The analogy thesis as just formulated needs at least two qualifications, however. The first is to recognize that core classification may occur with abstract as well as concrete terms. Take the core classification 'This is an animal', for instance. Its epistemic base might well be a concrete core classification that this is a paramecium and a linguistic truth that all paramecia are animals. If this is so, however, the concrete classification would be carried out by analogy.

Hence ‘This is an animal’ would be grounded mediately by analogy, not immediately as in concrete classifications. To cover the abstract case, I will claim that positive core classification is ultimately by analogy; negative core classification is ultimately by disanalogy.

A second qualification is needed for a relatively infrequent but key occurrence of core classification that is not analogical. When someone inaugurates a class term, there is no identification of the new by analogy with the old for the simple reason that there is no old. In the case of a newly-identified species, for example, there may be analogies between the first known individual of that species and members of neighboring species, but there can be no analogy within the species while only one exemplar is known. Call these seminal classifications coinages. The analogy thesis can then be restated: Except for coinages, positive core classification is ultimately by analogy; negative core classification is ultimately by disanalogy.

For all its brevity, the analogy thesis covers a lot of territory. The foregoing examples suggest that the thesis applies to physical objects, but the thesis in its most general form treats the core classification of events. This has immediate implications for physical objects, however, for each can be understood as an event that unfolds as long as the object exists. In addition, the thesis applies to actions, for actions are purposed events.

Understanding the analogy thesis in its full generality requires recognizing that core classifications occur in two ways. They can be linguistic, of course; someone says ‘That is a hake’, for example. But they can also be nonlinguistic, as in a perception that something is wet.[1] According to the analogy thesis, both types of core classification are effected ultimately by analogy. Hence analogy can be linguistic or not; we can state an analogy, or just perceive one.

3. Core Classification in Ethics

The preceding is all quite general; it makes no reference to specific domains of discourse. From here we could branch off into any domain at all. However, what I propose to explore is moral discourse of a specific kind. Loci for the sort of thing I have in mind are the early Platonic dialogues, which can be mined for insights on core classification in ethics: courage in the Laches, justice in Republic, Book I, temperance in the Charmides, and so forth.

Had the search for definitions in the early dialogues been successful, or had the definitions in Republic, Book IV, been more than rough-cut, stopgap measures, or had there been breakthroughs in the definition of ethical terms between Plato's time and our own, we could understand ethical core classification as follows. Imagine that we have an accurate definition of justice. It tells us that an action is just if, and only if, it is F and G and H. We could then classify individual actions by using our definition as a criterion: since this action is F and G and H, it is just; and that action, because it is F and G but not H, is not just.

There is ample reason by now to think that this is barking up the wrong tree. An extensive, multidisciplinary literature points toward a very different understanding of human classification.[2] Though I will not survey this literature here, I will note that Wittgenstein's remarks on family resemblance (1953, §65–78) are a point of departure for much later work in the field. And I will acknowledge the special importance of empirical work by Eleanor Rosch and her associates in cognitive psychology (Rosch 1973, 1978, 1983; Rosch and Mervis 1975; Rosch et al. 1976).[3]

Rosch is responsible for drawing together a number of separate empirical studies under the rubric of prototype theory. Part of the interest of prototype theory lies in its direct opposition to the classical conception of classification presupposed by Plato and passed on to scores of generations of western scholars. According to this classical view, class membership is determined by necessary and sufficient conditions, and accurate definitions state these

conditions. If this were the case, however, there would be no best examples of a kind; any member of the class would serve equally well, for all satisfy the same set of conditions. That there are best examples of a kind—prototypes—is a crucial result of Rosch's work. Speakers of American English, for example, consistently rate robins as better examples of bird than ostriches or penguins, and desk chairs as better examples of chair than beanbag chairs or electric chairs (Rosch et al. 1976). Class membership, so understood, is not an either-or affair; it is a matter of degree.[4]

Moral classes show prototype effects. We have no problem identifying Socrates' saving of Alcibiades' life as courageous, but how do we classify the suicide of Seneca's barbarian, who asphyxiated himself with the sponge he was given for wiping himself before his scheduled appearance in the circus to fight wild beasts? Even if we concur with Seneca's classification of it as courageous (1996, LXX.20–21, 66–69), it is not obviously so. This graded sort of membership in moral classes is due in part to moral education, which in its early stages proceeds through introducing prototypically moral and immoral actions in fairy tales and other narratives. But it is also an effect of what happens next. Once we have learned to manage a handful of moral predicates in prototypical situations, we begin to extend these terms to new situations by analogy.

I want to consider one example of this process in some detail. The example is trivial, in a sense, but that is precisely its point. It is a prototype, the kind of action that serves as a moral reference point for a community—primarily, in this case, that of the United States. (Non-American readers: forbearance, please. Prototypes are culturally embedded, and I have simply chosen one familiar to me.) The source for the incident is a classic biography of the American President George Washington (Weems 1962).[5]

Mason Weems was an Episcopal clergyman and bookseller who, a month after Washington's death in 1799, wrote to a business associate to propose a biography of Washington. His plan was to demonstrate that Washington's "unparrelled [sic] rise & elevation were due to his Great Virtues" (1962, xv). One of the virtues Weems attributed to Washington was honesty. To illustrate the point, he recounted the following incident:

> When George . . . was about six years old, he was made the wealthy master of a hatchet! of which, like most little boys, he was immoderately fond, and was constantly going about chopping every thing that came in his way. One day, in the garden, where he often amused himself hacking his mother's pea-sticks, he unluckily tried the edge of his hatchet on the body of a beautiful young English cherry-tree, which he barked so terribly, that I don't believe the tree ever got the better of it. The next morning the old gentleman finding out what had befallen his tree, which, by the by, was a great favourite, came into the house, and with much warmth asked for the mischievous author, declaring at the same time, that he would not have taken five guineas for his tree. Nobody could tell him any thing about it. Presently George and his hatchet made their appearance. George, said his father, do you know who killed that beautiful little cherry-tree yonder in the garden? This was a tough question; and George staggered under it for a moment; but quickly recovered himself: and looking at his father, with the sweet face of youth brightened with the inexpressible charm of all-conquering truth, he bravely cried out, "I can't tell a lie, Pa; you know I can't tell a lie. I did cut it with my hatchet." (1962, 12, Weems' emphasis)

The first point to be made about this story is that even though Weems attributes it to "an aged lady, who was a distant relative [of Washington], and when a girl spent much of her time in

the family," historians almost universally reject it as apocryphal (1962, 9, xxiv–xxxiv). The second point is the mordant one that Weems succeeded, apparently through falsehood, in placing what he took to be Washington's honesty before the "admiring eyes" of many children. By the time of Weems' death in 1825, twenty-nine editions of his Life had appeared; by 1925, the number had grown to eighty (1962, xx). Abraham Lincoln read it in "the earliest days of my being able to read," but so did many others (1962, xxii). The story of the cherry tree in particular reached millions through being excerpted in a vast body of Sunday-school books and textbooks, notably McGuffey's Readers, 120 million of which were published in the United States between 1836 and 1920 (1962, xx–xxiii, xlvi–xlviii; Ong 1982, 115–16). Given such a central place in moral education, there is little doubt that the incident has served as a base for Americans' understanding of honesty.[6] Almost in spite of itself, then, the incident became a moral prototype.

To see how such a prototype might be used in moral reasoning, let us return to Weems' account. From it we can extract an abstract description of the situation that includes the following features:

i) g is the child of f;

ii) g believes that p;

iii) g has a selfish desire that f not come to believe that p;

iv) f asks g whether p is true;

v) g conveys to f that p is true.

The result is a kind of template that is plainly applicable to other cases.

Let us consider a few. Incorrigible George next cuts down his father's apple tree, responding as before to his father's question. Because this recidivist action is identical to the prototype in the ways picked out by the template, the present action is honest as well. Now take an action that is like the preceding except that feature i) of the template is absent: f and g are not

related as parent to child. Despite the difference, the strong similarity between the prototype and this case would naturally lead us to classify the action as honest. A third case is like the second except that feature iii) of the template is missing as well (g is indifferent whether f comes to believe p or not, say). While the similarities between this case and the prototype make it easy enough to call the former honest, the exemplary honesty that would appeal to someone like Weems has been lost.

Our reactions to these increasingly divergent cases suggest that, as a matter of psychological fact, moral predicates are applied to novel actions or not on the basis of perceived similarities and dissimilarities between the actions and prototypes. This is most easily confirmed for predicates linked with "thick" concepts like honesty, brutality, and courage (Williams 1985, 129f). However, I submit that "thin" concepts such as right and good are tied to concepts like honesty via meaning postulates such as 'Honesty is right'. That is, thin core classifications like 'This action is right' are epistemically grounded in thick core classifications such as 'This action is honest', and these thick classifications are carried out via analogy. If this is so, then ethical core classification proceeds ultimately by analogy and disanalogy, and the analogy thesis holds for moral discourse in particular.

4. Norms for Analogies

The analogy thesis brings epistemological problems in its train. Suppose that perceptual analogies clash—one person perceives a color as mauve, say, while another perceives that it is not. Stating the conflicting analogies offers a way out by opening the analogies up to intersubjective criticism. But that is to lean on a slender and suspect base: the much-maligned argument by analogy.

This section of the paper addresses the argument by analogy in two stages. It identifies a few of the many varieties of argument by analogy, first of all. It then addresses the key question of how to distinguish good arguments by analogy from bad ones.

J. S. Mill once remarked that "There is no word . . . which is used more loosely, or in a greater variety of senses, than Analogy" (1973–74, III.20.1, 554). So let us rehearse some distinctions among kinds of analogy. Although the following typology refers exclusively to linguistic analogy, it should be kept in mind that nonlinguistic analogies can be described as if they were linguistic. That is, a perceived analogy between two things can be described as if an analogical argument concerning them were carried through.

One place to begin a typology of analogy is with the conspicuous divide between general analogies, analogies with at least one quantified sentence, and singular analogies, composed entirely of unquantified sentences. The latter, which are indispensable to core classification, can be subdivided into perfect and imperfect forms. In perfect analogies the relata are believed to share all relevant properties. A simple example is the traditional argument by analogy:

A_1: Fa & Ga.
Fb.
Hence Gb.

In imperfect analogies, on the other hand, the relata are believed to share some but not all relevant properties. One example from the literature (Pietarinen 1972, 68–69) is

A_2: Fa & Ga.
–Fb.
Hence Gb.

A_2 makes the imperfectly analogical claim that a and b share the property G but not the property F. Though the perfect-imperfect distinction is just the beginning of a typology of singular analogy, we need pursue the matter no further here.[7]

Instead, let us turn to the crucial normative question: What is the difference between good and bad analogies? To sketch an answer, suppose we pull out the old critical saw about sound argumentation. We recall that for an argument to be deductively sound it must meet at least two necessary conditions. A condition on the argument's content requires all of its premises to be true. And a condition on the argument's form requires it to be valid in the sense that it is impossible that its conclusion is false when its premises are true.

Along the same lines, we can specify two necessary conditions for an argument to be inductively cogent. The content condition remains the same: all the argument's premises must be true. But the formal condition is different; it must be weaker than deductive validity yet still demanding. Here the usual requirement is inductive strength, which stipulates that it be improbable that the argument's conclusion is false when its premises are true (Skyrms 1986, 7). Now it is improbable that the conclusion is false when the premises are true if, and only if, it is probable that the conclusion is true when the premises are true. For an argument to count as inductively strong, then, the conditional probability of its conclusion given the premises must be greater than or equal to that of any rival conclusion based on the same premises.

Hence the proposed standard of inductive cogency amounts to this: An argument is inductively cogent only if

a) all the argument's premises are true; and

b) the conditional probability of the argument's conclusion given its premises must be greater than or equal to that of any rival conclusion based on the same premises.

Before relating these inductive conditions to arguments by analogy, let us note how neatly they dovetail with our deductive practice. The condition on deductive content is exactly the same, as we have noted. The condition on deductive form, the requirement that the argument be structured such that if the premises are true, then the conclusion must be true, actually implies the inductive condition on form. That is, if an argument is deductively valid, then it is also inductively strong, for its conclusion has a greater probability on the premises (probability 1) than any rival conclusion based on the same premises (probability 0). The condition on deductive form is thus a special case of the inductive condition on form. This was Wittgenstein's point in the Tractatus: "The certainty of logical inference is a limiting case of probability" (Wittgenstein 1922, 5.152; cf. Haack 1978, 17).

Now let us link the foregoing to analogy. No argument by analogy is deductively sound, but they are not all equally unsound. To distinguish the better from the worse, I propose that we treat argument by analogy as one form of inductive argument. This is a time-honored view. Mill, for instance, remarked that arguments by analogy are "supposed to be of an inductive nature" (1973–74, III.20.1, 554), and Carnap handled analogy as induction from at least 1945 on (1945, 87–88). Adopting this proposal would permit us to distinguish good and bad arguments by analogy by applying the standard for inductive cogency outlined two paragraphs above.

Before working out the implications of this proposal, I want to address a possible objection. It might be admitted that many analogies are inductive while still maintaining that many are not; many are abductive, one might claim, rather than inductive. In response, I suggest that we distinguish logical and functional views of argumentation. The logical approach is to classify an argument according to the degree of support the conclusion receives from the premises. But the functional approach keys on how the argument is used. Peirce's trichotomy of deductive, inductive, and abductive arguments is (usually) functional, for instance. Abduction is

the first step of scientific reasoning; it advances a hypothesis. Induction is the last step; it uses experiment to verify a deductive consequence of the hypothesis (Peirce 1931–1958, 2.96, 7.218). One Peircean example treats "This is an ex-priest" first as the conclusion of an abductive argument to explain a surprising conjunction of features and then, after the "experiment" of getting the man to remove his hat to confirm that he is tonsured, as the conclusion of an inductive argument (Peirce 1902).

The functional difference between abductive and inductive arguments is large, but the logical difference, I submit, is small. That this man is an ex-priest follows with some probability from premises describing the initial set of features (Peirce does not name them), and it follows with some greater probability from the initial set plus the premise on tonsure. The difference is one of degree, as Peirce seems to recognize (Peirce 1931–1958, 2.624). But since both conclusions have a conditional probability less than 1, both arguments contrast with demonstrative arguments, whose conclusions have a conditional probability of 1. Both arguments are therefore inductive in the standard logical sense of being nondemonstrative (e.g., Carnap 1945, 72; Skyrms 1986, 6–7; Adams 1998, 70). We can recognize that an argument is functionally abductive, then, and at the same time logically inductive. The foregoing claim that analogy is inductive should be understood in this logical sense.

Arguments by analogy offer difficulties, but no peculiarly analogical difficulties, in determining the truth of their premises. Yet they do present special difficulties over form. How might we go about applying the condition of inductive strength? Arguments by analogy are built around the relation of similarity, so intuitively it would seem that the relata of inductively strong analogies are somehow more similar than dissimilar, while those of inductively weak analogies are somehow more dissimilar than similar. But putting this intuition to work would require some sort of similarity metric. Where could we find one?

I propose that we consider those logics developed along lines sketched out by Wittgenstein (1922, 5.15–5.156) and Waismann (1930–31). Carnap (1952) made the decisive step forward, and his work has served as the basis for later advances by Hintikka (1966), Carnap (1971, 1980), Pietarinen (1972), Hintikka and Niiniluoto (1976), Kuipers (1978, 1984), Niiniluoto (1981), Skyrms (1991, 1993), and Festa (1997), among others. In Carnap's mature work (e.g., 1942, 96–97; 1945, 73–75), the concept of range is a semantic concept explicable as the set of models in which a given sentence (or conjunction of sentences) is true. Suppose we call such models the sentence's alethic models.

The relationship between the alethic models of an argument's premises and those of its conclusion shows the probability of the conclusion on the evidence of the premises. There are two types of cases. If the alethic models of the conclusion include all the alethic models of the premises, the conclusion follows from the premises with probability 1, and the argument is deductively valid. On the other hand, if the alethic models of the conclusion do not include all of the alethic models of the premises, the conclusion follows with some probability less than 1, and the argument is not deductively valid. For example, if 3/4 of the premises' alethic models are included in those of the conclusion, the probability of the conclusion given the premises is 3/4.

One result of Carnap's critique of classical probability was his λ-continuum of inductive methods (1952). Given any method of this continuum, the degree of confirmation of a singular hypothesis on the evidence lies within an interval bounded by an empirical factor and a logical factor. The empirical factor is the evidence e_Q, the ratio of the n_Q favorable instances of some strongest property Q to the total number n of instances examined. The logical factor is equal to relative width, which is very roughly the coverage of a property relative to the totality of properties the language admits. Less roughly, for a first-order language with identity recognizing a finite number m of logically independent primitive properties, there are $2^m = K$ strongest

properties in the language. Any property that can be picked out in the language is either a strongest property or equivalent to a disjunction of strongest properties. If the property is a strongest property, its relative width is 1/$\underline{K}$. If not, being equivalent to a disjunction of $\underline{w}$ strongest properties instead, its relative width is $\underline{w}$/$\underline{K}$.

Exactly what point of this interval represents degree of confirmation—or probability, as I shall say—is determined by taking a weighted mean of the empirical and logical factors. Different λ-methods use different logical weights, that is, different specifications of the parameter λ, which can take values from 0 to ∞ inclusive. Now suppose that we have evidence $\underline{e}_Q$ and that λ can vary with $\underline{K}$ but not with $\underline{n}_Q$ and $\underline{n}$. Then, for any method of the continuum, the

$$p(h_Q, e_Q) = \frac{n_Q + \frac{\lambda(K)}{K}}{n + \lambda(K)}. \tag{1}$$

probability $\underline{p}$ of the hypothesis $\underline{h}_Q$ that the next individual will have a strongest property is (1) allows for uncountably many λ-methods, but Carnap's favorite was c*, where $\lambda(\underline{K}) = \underline{K}$. For consistency with our probabilistic terminology, I will refer to this method as 'p*' and express its representative function as[8]

$$p^*(h_Q, e_Q) = \frac{n_Q + 1}{n + K}. \tag{2}$$

The methods of the λ-continuum are problematic in several ways, but the crucial shortcoming for our purposes is their handling of analogy. For example, where $\underline{K} = 4$, p* assigns the perfect analogy A_1 a probability of 2/3, which seems reasonable enough, while the imperfect analogy A_2 receives a probability of 1/2, which also seems reasonable enough—until we notice

that its property analogy has been completely overlooked. That is, since A_2's conclusion '$\underline{Gb}$' has a probability of 1/2, the other possible conclusion, '–$\underline{Gb}$', receives the same probability. But that is to consider the conjunction of the disanalogous properties $\underline{FG}$ and $\underline{\overline{F}\overline{G}}$ just as likely as that of A_2's analogous properties $\underline{FG}$ and $\underline{\overline{F}G}$.

The λ-continuum has been superseded by Hintikka's α-λ continuum (1966), which extends the λ-continuum to improve the handling of inductive generalization; Carnap's Basic System (1971, 1980), which extends the λ-continuum by including predicates of unequal as well as equal widths; and Hintikka's and Niiniluoto's $\underline{K}$-dimensional system (1976), which axiomatizes a substantial portion of the α-λ continuum.[9] But all these systems have the same difficulty with singular analogy (Welch 1999). Various remedies have been proposed.[10] That to be pressed into service here originated with Kuipers (1984) as a counter-proposal to one by Niiniluoto (1980, 1981), who subsequently endorsed it (1988, 287).

Kuipers observes that we can view (1) as the application of the straight rule to $\underline{n}_Q$ real empirical instances of the strongest property $\underline{Q}$ and $\lambda(\underline{K})/\underline{K}$ virtual logical instances of $\underline{Q}$ (1984, 68–69). Why not then account for analogy by analogy with these virtual logical instances? That is, why not add virtual analogical instances of $\underline{Q}$ to factor in the relative similarities of properties?[11] Let each strongest property $\underline{Q}$ be associated with $\alpha_Q(\underline{e})$ virtual analogical instances that represent $\underline{Q}$'s similarity to the properties of the evidence. When $\underline{n} \geq 1$, $\alpha_Q(\underline{e})$ is > 0, but when $n = 0$, the absence of evidence requires that $\alpha_Q(\underline{e}) = 0$. In addition, let $\alpha(\underline{e})$ virtual analogical instances represent the summation of similarities that all strongest properties have to the properties of the evidence. The ratio $\alpha_Q(\underline{e})/\alpha(\underline{e})$ would then indicate $\underline{Q}$'s portion of total similarity to the evidence. This ratio could therefore be added to (1) as an analogy factor

comparable to the empirical and logical factors. Like the empirical and logical factors, the various analogy factors sum to 1. Where $0 < \lambda < \infty$, (1) would become:[12]

$$(3) \qquad p(h_Q, e_Q) = \frac{n_Q + \frac{\lambda(K)}{K} + \alpha_Q(e)}{n + \lambda(K) + \alpha(n)} .$$

Accordingly, the representative function for p* would be adapted for the new method p**:

$$(4) \qquad p^{**}(h_Q, e_Q) = \frac{n_Q + 1 + \alpha_Q(e)}{n + K + \alpha(n)} .$$

Since there are an unlimited number of ways of assigning these numbers, how could we determine the appropriate number of virtual analogical instances? Niiniluoto has described a natural way of measuring degrees of resemblance among strongest properties (Niiniluoto 1981, 12–14).[13] Where $\underline{d}_{uv}$ is the number of primitive properties not shared by the strongest properties $\underline{Q}_u$ and $\underline{Q}_v$, their degree of resemblance $\underline{r}$ can be expressed as

$$(5) \qquad r_{uv} = \frac{1}{1 + d_{uv}} .$$

Given primitive properties $\underline{F}$ and $\underline{G}$, for example, (5) determines the degrees of resemblance between the strongest property $\underline{FG}$ on the one hand and $\underline{FG}$, $\underline{F\overline{G}}$, $\underline{\overline{F}G}$, and $\underline{\overline{F}\overline{G}}$ on the other to be 1, 1/2, 1/2, and 1/3 respectively.

(5) affords a particularly simple way of determining appropriate analogy factors. Suppose initially that the evidence manifests just one strongest property. Where $\underline{Q}_u$ is this strongest

property and $\underline{Q}_v$ is the strongest property of the hypothesis $\underline{h}_Q$, let the value of $\underline{r}_{uv}$ be $\alpha_Q(\underline{e})$. Then, where $\underline{Q}_u$ is once again the strongest property of the evidence, total similarity $\alpha(\underline{e})$ would be

(6) $$\sum_{v=1}^{K} r_{uv}.$$

If the strongest property of the evidence is $\underline{FG}$, then the analogy factors for $\underline{FG}$, $\underline{F\overline{G}}$, $\underline{\overline{F}G}$, and $\underline{\overline{F}\overline{G}}$ would be $\frac{1}{7/3}$, $\frac{1/2}{7/3}$, $\frac{1/2}{7/3}$, and $\frac{1/3}{7/3}$ respectively. The factors are expressed in unreduced form to highlight the conceptual links with (5) and (6).

In more complicated cases where more than one strongest property appears in the evidence, $\alpha_Q(\underline{e})$ is just the sum of the values of $\underline{r}_{uv}$ for each property $\underline{Q}_u$ of the evidence and the property $\underline{Q}_v$ of the hypothesis $\underline{h}_Q$. For $\alpha(\underline{e})$ we note that the value of (6) for any strongest property of the evidence equals the value of (6) for any other strongest property of the evidence, though the individual values of $\underline{r}$ are distributed differently. Hence where the number of strongest properties instantiated by the evidence is $\underline{i}$, $\alpha(\underline{e})$ is generally

(7) $$i\sum_{v=1}^{K} r_{uv}.$$

Examples appear in Section 5 below.

Applying p** along these lines reflects the property analogies that unmodified p* does not. As we have seen, p* with $\underline{K} = 4$ allots probabilities of 1/2 to both A_2's more similar conclusion '$\underline{Gb}$' and the less similar conclusion '$\underline{\text{–}Gb}$'. Under the same assumptions, however, p** with analogy factors of $\frac{1/2}{7/3}$ for $\underline{\overline{F}G}$ and $\frac{1/3}{7/3}$ for $\underline{\overline{F}\overline{G}}$ assigns probabilities of 9/17 (about .53) to A_2's more similar conclusion '$\underline{Gb}$' and 8/17 (about .47) to the less similar conclusion '$\underline{\text{–}Gb}$'. This is not an isolated instance. p** is sensitive to property analogy wherever p* is not. We

can use it, therefore, to estimate the probability on the premises of any singular analogical conclusion whatever.[14]

Although I have limited myself to p* for ease of illustration, any of an infinite number of alternative methods can be property-sensitized in the same way. Yet we brush up against a well-known difficulty in doing so: there are, after all, so many of these methods. Since different methods give different values, how do we know which one to choose? This is indeed a problem, but it seems not to have been noticed that there are situations where this embarrassment of methods does not matter at all. The reason is this: knowing merely that one conclusion is more probable than its rivals is sometimes enough; exactly how much more may be superfluous information. Suppose that to be the case with A_2, for example. Yet even though the probabilities of '$\underline{Gb}$' and '$-\underline{Gb}$' on A_2's premises vary from method to method, their comparative relations do not. '$\underline{Gb}$' in this context is always more probable than its rival '$-\underline{Gb}$'.

In these cases, probability can profitably be compared to temperature. There are alternative temperature scales, but since jumping from one to another preserves relations of hotter than and colder than, the main thing is to pick a scale and stick with it. Similarly, if all we need to know is which conclusion is more probable, the choice of a method is relatively unimportant. One such situation is described in the following section.

To conclude this section, let us consider the bearing of these methods on non-monotonic reasoning. Take the stock argument about Tweety, by now something of a celebrity in this area:

Birds can fly.

Tweety is a bird.

Thus Tweety can fly.

Now if we add the premise ‘Tweety is a penguin’ to the original premises, we get the conclusion ‘Tweety cannot fly’. The conclusions of the initial argument (call it Inference 1) and the augmented argument (Inference 2) are plainly contradictory.

This example appears to cover three basic cases. 1) The first premise might mean ‘All birds can fly’. If so, this premise would be false, and we would reject both Inference 1 and Inference 2. 2) The first premise might mean ‘Most birds can fly’, and we might know that Tweety is a penguin in drawing Inference 1. Then Inference 1 would violate the requirement of total evidence, and we would therefore reject it. 3) The first premise might mean ‘Most birds can fly’, and we might not know that Tweety is a penguin in drawing Inference 1.

Case 3) covers three subsidiary cases. 3a) When drawing Inference 1, we might know that penguinhood is relevant but just not know how to classify Tweety. Then, when we learn that Tweety is a penguin, we would reject Inference 1. 3b) When drawing Inference 1, we might know that penguinhood is relevant but falsely believe that Tweety is not a penguin. Then, when we discover that Tweety is a penguin, we would reject Inference 1. 3c) When drawing Inference 1, we might not know that penguinhood is relevant. Then, when we discover that this background assumption is false, we would reject Inference 1.

Of these multiple possibilities, 1), 3b), and 3c) are quickly decided by appeal to the truth condition on cogent argumentation, and 2) just as quickly by the requirement of total evidence. The remaining case, 3a), is the only one directly relevant to the inductive methods under discussion. The application, I suggest, should be as follows. The conclusions of both Inference 1 and Inference 2 should be understood as following with some probability from their respective premises. The representative function for the method of choice, p** for example, could provide the probability that Tweety can fly given that we do not know whether or not she is a penguin (Inference 1). It could also determine the probability that she can fly given that we know that she

is a penguin, and hence the probability that she cannot fly based on the same evidence (Inference 2). The problem is basically one of updating probabilities, and the representative function would yield the prior and posterior values for the probability transition.

5. Norms for Ethical Analogies

We alluded to the early Platonic dialogues in section 3, and we must briefly return. Despite their abstract definitional concerns, these dialogues are rooted in a practical problem that is never far beneath the surface. The Euthyphro is particularly explicit. In the dialogue's discussion of piety, Socrates identifies two types of disagreement: those that cause hatred and social discord, and those that do not (1975, 7b–d, 9). Those that do not include differences over number, size, and weight; those that do concern "the just and the unjust, the beautiful and the ugly, the good and the bad" (1975, 7d, 9).[15] These socially divisive disagreements are predominantly moral. However, Socrates gets Euthyphro to see that people do not disagree about moral questions as such; they agree that the wrongdoer should be punished, for example, though they may disagree about whether someone is a wrongdoer (1975, 8c–d, 11). These disagreements are disputes "about each action. . . . Some say it is done justly, others unjustly" (1975, 8e, 11).

These socially divisive disagreements have a plainly identifiable root: the vagueness of terms like 'just'. The Euthyphro is built upon just this sort of occurrence: Euthyphro says his prosecution of his father is pious, but his family says it is not; Socrates' friends claim that Socrates' actions are pious, but Meletus counters that they are not. The term 'pious' is evidently vague. Structurally similar disagreements over whether an individual action is just or courageous or honest are just as troublesome in our own time as they were in Plato's, and they stem just as clearly from the vagueness of moral terms.

I will refer to disagreements of this concrete and socially divisive sort as clash points. Given their attendant social problems, clash points raise urgent normative questions: In such cases, is there a rational way of choosing sides? Provided the interlocutors are willing to give reasons for their views, there is. According to the analogy thesis, reasons for these clashing classifications must ultimately be arguments by analogy, and, as we have seen in the last section, good analogies can be distinguished in principle from bad ones. If either of the analogies has a false premise, the analogy should be rejected. But if there is agreement over the truth of the premises, the disagreement over the conclusion must be rooted in conflicting views of the similarity between the controversial act and its prototype(s). That is a resolvable disagreement. Similarity and dissimilarity come in degrees, and quantitative inductive logics like p** provide metrics for measuring them.

The procedure is to form what Kemeny (1963, 722) has called the "minimal language," the simplest language containing the singular and general terms of both premises and conclusion, and determine the conditional probability of the conclusion on the premises with the help of the representative function. If one of the analogies has the form of A_2, for instance, where a is the prototype, b is the disputed act, and G is the controverted moral property, we could argue that due to measurable degrees of similarity, we have more reason than not to assign G to b (details in Section 4). The predicate 'G' would continue to be vague, but it would no longer be vague at the clash point.

Suppose we try this out on an example that is complex enough to model real-world difficulties. In De Officiis, Cicero presents the case of the grain merchant:

> . . . if, for example, an honest merchant has brought a great quantity of grain from Alexandria to Rhodes at a time when the Rhodians are suffering from great famine and the price of grain is high, and if he knows that more merchants have set sail from

> Alexandria and has seen their ships on his way sailing in the direction of Rhodes laden with grain, is he to tell the Rhodians or keep quiet and get the best price he can for his cargo? We can assume that he is a wise and honest man, and can for the purposes of our discussion take it that he would not conceal the fact from the Rhodians if he thought it dishonest, but he is in doubt about its honesty. (1967, III.12.50, 153–54)

Cicero indicates that the case was a staple of Stoic moral discourse, and that it was a point of contention between Diogenes of Babylon and his disciple Antipater. Diogenes argued that since the grain merchant has not been asked whether the other ships are on the way, to say nothing about them is consistent with honesty. Antipater, on the other hand, thought silence dishonest. Diogenes and Antipater have gone the way of all flesh, but their disagreement has not; it can be revived in almost any contemporary audience. The term 'honest' is vague, and what we have is a genuine clash point.

To have any prospects for resolving it rationally, the disputants must agree about three things. There must be some consensus on prototypes, first of all, on certain actions as bearers of the problem predicate. This is no special requirement, however; agreement on the premises is indispensable for reaching consensus through any kind of inference, inductive or deductive. So let us assume that, as is sometimes the case, the disputants are looking back at the case from a later point in time—our own, for example. Let us say that while they have several disagreements about honesty, there is one and (to keep it simple) only one point of agreement: the Weemsian prototype. They agree that it was honest.

Just as important as agreement on a clear positive instance of the predicate is agreement on a clear negative.[16] Having agreed on the Weemsian prototype, we might suppose that George is asked whether he cut down the cherry tree but that he answers differently: "I can't tell a lie, Pa; you know I can't tell a lie. I did <u>not</u> cut it down with my hatchet." This action is clearly

dishonest, and our knowledge that it is so should be included as a premise of the argument.

The final matter for agreement is what features of the case are morally relevant. Like all induction, moral induction requires selecting out, from all the properties of the objects under discussion, those that are relevant to the question.[17] These are the only properties that figure in the minimal language, and they are the primitive properties that make up the K strongest properties for determining relative width. Suppose, then, that the disputants agree that, with one exception, the template of Section 3 identifies the morally relevant features. The exception is feature i), the filial relation, which is counted out as irrelevant. Consequently, the features agreed to be relevant are: g's belief that p (call this property B for short); g's selfish desire that f not come to believe p (property D); f's asking g whether p is true (property A); and g's conveying to f that p is true (property C).

Where property H is g's honesty, a is the Weemsian prototype, b is the negative instance where George is asked but conveys misinformation, and c is the problem case of the merchant, the analogical argument leading to Antipater's conclusion that the grain merchant's silence is not honest can be represented as:

A_3: Aa & Ba & Ca & Da & Ha.

Ab & Bb & –Cb & Db & –Hb.

–Ac & Bc & –Cc & Dc.

Hence –Hc.

Once the argument is explicitly set out, determining whether 'Hc' or '–Hc' is the better conclusion is straightforward. The question of whether the grain merchant's case is more similar to the clear positive or the clear negative case is actually evident by inspection:

Positive: ABCD

Grain merchant: $\bar{A}B\bar{C}D$

Negative: $AB\bar{C}D$

In fact, the greater similarity of the merchant's case to the clear negative is easily quantified. Availing ourselves of (5), Niiniluoto's measure of resemblance between strongest properties, we have

$r(ABCD, \bar{A}B\bar{C}D) = 1/3$

$r(AB\bar{C}D, \bar{A}B\bar{C}D) = 1/2.$

The process of quantification can be carried still a step further by recalling the methods of Section 4. Because the disputants believe the argument's premises to be true, the clash over the conclusion must stem from divergent estimates of the support provided by the premises to the conclusion. Suppose we use p** to mediate. Since there are five primitive properties, there are 2^5 = 32 strongest properties. The premises assert the strongest properties $\underline{ABCDH}$ and $\underline{AB\bar{C}D\bar{H}}$, and the rival conclusions (together with the third premise) imply the strongest properties $\underline{\bar{A}B\bar{C}DH}$ and $\underline{\bar{A}B\bar{C}D\bar{H}}$. To represent the relevant similarities, we can rely on (5) and (7). They determine analogy factors of $\frac{1/4}{21} + \frac{1/2}{21} = \frac{3/4}{21}$ for $\underline{\bar{A}B\bar{C}D\bar{H}}$ and $\frac{1/3}{21} + \frac{1/3}{21} = \frac{2/3}{21}$ for $\underline{\bar{A}B\bar{C}DH}$. Then p** assigns a value of 21/41 (about .512) as the probability of A_3's conclusion on its premises and a probability of 20/41 (about .488) to the rival conclusion that the merchant's silence is honest. Marginally, then, and relative to the Washingtonian prototypes, the grain merchant's silence is more dishonest than honest. The extension of 'honest' could therefore be clarified by subtracting the action of the grain merchant's silence. This would be one pass in the inductive molding of the concept of honesty.

Whatever else we may want to say about this conclusion, I think it accords well with at

least one of our convictions about the case. Since the rival conclusions are intuitively very close, an assignment of .80, say, to one of them would clearly have been wrong. Nevertheless, this conclusion may not jibe with all of our convictions. It should do so only if our convictions are based on exactly the same evidence. The conclusion that the grain merchant's silence is honest can be shown to follow from a different evidential base. But it does not follow from this one.

As we noted above in connection with A_2, any of p**'s property-sensitized cousins will give slightly different probabilities for A_3's conclusion. Knowing which are the exact values would indeed be welcome. But here we are fortunate; that is beside the point. To resolve disagreements over clash points like the grain merchant's, all we really need to know is which conclusion is more probable, and that has been accomplished. It is an elementary matter to show that the greater probability of '–Hc' relative to 'Hc' on the premises of A_3 is invariant across these methods.

Up to this point our discussion of the grain merchant has tacitly assumed that the relevant predicates have equal logical weight. This assumption need not always hold, and it can be jettisoned at will. The key is to take a hint from Carnap's Basic System, which accommodates unequal as well as equal logical weights (Carnap 1971, 1980). Each strongest property Q can be correlated with $\lambda_Q > 0$ virtual logical instances. The sum of the various λ_Q is λ, which represents total logical weight. Hence each strongest property can be outfitted with a logical factor λ_Q/λ that expresses its share of the system's logical weight. Like empirical and analogy factors, these logical factors must sum to 1. The result is a generalization of (3), which is Kuipers' replacement for the representative function of the λ-continuum:

(8) $$p(h_Q, e_Q) = \frac{n_Q + \lambda_Q + \alpha_Q(e)}{n + \lambda + \alpha(e)} .$$

In the case of the grain merchant, for example, we might regard any strongest property that includes the combinations $\underline{\text{ACH}}$ or $\underline{\text{A}\overline{\text{C}}\overline{\text{H}}}$ as particularly significant. We might be willing to allot more logical weight to the 8 strongest properties in which these combinations appear than to the remaining 24 properties. Instead of the single virtual logical instance employed under p**'s assumption of equal logical weight, we might apportion 1.5 logical instances to the more significant properties and .5 instances to the less significant ones. Since 8 "heavy" properties and 24 "light" ones would require a total of 24 virtual logical instances, each heavy property would have a logical factor of $\frac{3/2}{24}$ and each light property a logical factor of $\frac{1/2}{24}$. Like the analogy factors, the logical factors are expressed in unreduced form in order to make the number of their virtual instances transparent.

Using these logical factors and the same analogy factors as before, (8) determines the probability of A_3's conclusion '–$\underline{\text{H}}$' on its premises to be 15/29 (about .517) and that of the rival conclusion '$\underline{\text{H}}$' to be 14/29 (about .483). Note that the probability of '–$\underline{\text{H}}$' is slightly higher when the strongest property $\underline{\overline{\text{A}}\text{B}\overline{\text{C}}\text{D}\overline{\text{H}}}$ receives .5 logical instances than when it receives 1 such instance (using p** above). Under our assumptions, (8) applied to A_3 is a decreasing function that approaches .5 as a limit when $\lambda_Q \rightarrow \infty$.

6. Conclusion

The key idea of Section 5 is to apply the inductive criterion for cogent arguments by analogy outlined in Section 4. As with all logical tools, however, the application can be arduous. First of all, as we have noted, consensus through inference cannot be achieved without prior

agreement on the premises. In the deductive case, vagueness, ambiguity, and lack of empirical information may create disagreement over the truth of an argument's premises, and the result can be disagreement over the soundness of the argument. The same snags can wreak the same havoc with inductive cogency. They can, in extreme cases, even impede consensus over prototypes. Second, inability to agree on which features are morally relevant can actually prevent the premises from being formulated in a mutually acceptable way. One person might formulate premises with one set of relevant features, and another might insist on a divergent set. Finally, just as the deductive validity of very complex arguments can be determined in principle but only with difficulty in practice, the same is true of inductive strength. The evidence for a given conclusion may be quite complex. Nevertheless, the inductive methods discussed in this paper are complete: for any noncontradictory premises and conclusions based on them that are formulable in the language for which the method is defined, the probability of the conclusion given the premises can be determined in principle (Carnap 1952, 16–18, 30–32).

Despite these difficulties of application, the criterion for cogent inductive arguments introduced in Section 4 can be plied to good effect. Section 5 uses this criterion to reduce vagueness in the case of the grain merchant. The focal predicate of this case is ethical, but the vagueness of nonethical predicates should be reducible in the same piecemeal way. Though this is a hypothesis subject to further investigation, the formality of the present approach already provides some confirmation. For the conclusion of any argument isomorphic to A_3 would be more probable than a contradictory conclusion based on the same premises, regardless of whether the constituent predicates are ethical or not. In addition, this inductive procedure has the systematic virtue of full coherence with our standard criterion for sound deductive arguments, making it a natural candidate for a principled approach to vague predicates.[18]

NOTES

[1] Typically, classification reflects belief. In my Humean sense of the term, beliefs can be linguistic or not; perceptions, for example, are nonlinguistic beliefs. This opens up lines of inquiry into the controversial arena of animal belief. Plato and Aristotle split on the matter, as Richard Sorabji (1993) interprets them, with Plato attributing beliefs to animals and Aristotle not. In contemporary philosophy, Jeffrey (1985) is Platonic on this point while Davidson (1982) is Aristotelian. But I will not pursue these matters here.

[2] Lakoff (1987, ch. 2) surveys much of this literature.

[3] Rosch (1978) is conveniently reprinted in Margolis and Laurence (1999, 189–206) along with several papers that aim to supplement or replace prototype theory.

[4] Even the classes recognized by the physical sciences seem generally to be graded. In discussing the theoretical identities 'Water is H_2O' and 'Temperature is mean kinetic energy', Putnam (1983, 64) remarks that "the 'essence' that physics discovers is better thought of as a sort of paradigm that other applications of the concept ('water', or 'temperature' must resemble than as a necessary and sufficient condition good in all possible worlds."

[5] The incident appears for the first time in the fifth edition of 1806. The citations from Weems (1962) are from the ninth edition of 1809.

[6] To avoid confusion over the kind of honesty that has to do with property, one might prefer the term 'truthfulness' here. I use 'honesty' to connect up with an example below.

[7] A detailed typology of singular analogies appears in Welch (1999, pp. 209–13).

[8] Representative functions are so-called because they determine all other values within the system. Carnap replaced the term 'characteristic function' of (1952) with 'representative function', and his later usage is followed here.

[9] On the relation between the α-λ continuum and the K-dimensional system, see Kuipers (1978, 262).

[10] For the α-λ continuum, see Hintikka (1968, 228; 1969, 28–33) and Pietarinen (1972, 91–99). For the K-dimensional system, see Niiniluoto (1980, 1981, 1988), Spohn (1981), Costantini (1983), and Kuipers (1984).

[11] My development of this idea differs somewhat from Kuipers'.

[12] The resulting systems are unusual in that they are not indifferent to the order in which predicates are instantiated, thereby violating the axiom of individual symmetry upheld by Carnap (1952, 14; 1963, 975) and others (e.g., Maher 2000, 64). Nevertheless, the probabilities obtained from the various orders of instantiating predicates all converge to the same point (Kuipers 1984, 76). For those unwilling to give up the axiom of symmetry, steps toward a satisfactory treatment of analogy may be found in the work of Skyrms (1993) and Festa (1997).

[13] Kuipers (1984, 67, 73–74) and Niiniluoto (1988, 279–280) offer alternative measures.

[14] Carnap (1963, 75, 973–974) came to view these methods as approximations, and they are so regarded here.

[15] The ensuing discussion builds on Welch (1997, pp. 1018–21).

[16] I am indebted to Jesse Hughes for spotting a difficulty in an earlier version of this argument that spurred the following paragraph.

[17] Though the issue of relevance deserves more attention than I can give it here, I expect both widespread agreement and occasional disagreement over what counts as morally relevant. Such disagreement, if it arises, need not be unresolvable; we have, I think, resources with which to work. But the point for the moment is that 'relevant' is a vague term. Just as disagreement over

borderline cases of a vague term may prevent determination of deductive soundness, disagreement over borderline cases of relevance may do the same for inductive cogency.

[18] I am grateful to Theo Kuipers and two anonymous reviewers for very helpful comments on an earlier version of this paper.

REFERENCES

Adams, E. (1998): A Primer of Probability Logic, Stanford: CSLI Publications.

Carnap, R. (1942): Introduction to Semantics, Cambridge, MA: Harvard University Press.

Carnap, R. (1945): 'On Inductive Logic', Philosophy of Science 12, 72–97.

Carnap, R. (1952): The Continuum of Inductive Methods, Chicago: University of Chicago Press.

Carnap, R. (1963): 'Probability and Inductive Logic' and 'My Basic Conceptions of Probability and Induction', in Schilpp (ed.), pp. 71–77, 966–979.

Carnap, R. (1971): 'A Basic System of Inductive Logic, Part 1', in R. Carnap and R. Jeffrey (eds.), Studies in Inductive Logic and Probability, Vol. I, Berkeley: University of California Press.

Carnap, R. (1980): 'A Basic System of Inductive Logic, Part 2', in R. Jeffrey (ed.), Studies in Inductive Logic and Probability, Vol. II, Berkeley: University of California Press.

Cicero (1967): On Moral Obligation: A New Translation of Cicero's De Officiis, trans. John Higginbotham, London: Faber. De Officiis originally published 45–44 B.C.

Costantini, D. (1983): 'Analogy by Similarity', Erkenntnis 20, 103–114.

Davidson, D. (1982): 'Rational Animals', Dialectica 36, 318–27. Rpt. in LePore and McLaughlin (eds.), pp. 473–480.

Festa, R. (1997): 'Analogy and Exchangeability in Predictive Inferences', Erkenntnis 45, 229–252.

Haack, S. (1978): Philosophy of Logics, Cambridge: Cambridge University Press.

Hintikka, J. (1966): 'A Two-Dimensional Continuum of Inductive Methods', in Hintikka and Suppes (eds.), pp. 113–132.

Hintikka, J. (1968): 'Induction by Enumeration and Induction by Elimination' and 'Reply', in I. Lakatos (ed.), The Problem of Inductive Logic, Amsterdam: North-Holland, pp. 191–216, 223–231.

Hintikka, J. (1969): 'Inductive Independence and the Paradoxes of Confirmation', in N. Rescher (ed.), Essays in Honor of Carl G. Hempel, Dordrecht: D. Reidel, pp. 24–46.

Hintikka, J. and I. Niiniluoto (1976): 'An Axiomatic Foundation for the Logic of Inductive Generalization', in M. Przełęcki, K. Szaniawski, and R. Wójcicki (eds.), Formal Methods in the Methodology of Empirical Sciences, Dordrecht: D. Reidel and Wrocław: Ossolineum, pp. 57–81.

Hintikka, J. and P. Suppes, eds. (1966): Aspects of Inductive Logic, Amsterdam: North-Holland.

Jeffrey, R. (1985): 'Animal Interpretation', in LePore and McLaughlin (eds.), pp. 481–487.

Kemeny, J. (1963): 'Carnap's Theory of Probability and Induction', in Schilpp (ed.), pp. 711–738.

Kuipers, T. (1978): 'On the Generalization of the Continuum of Inductive Methods to Universal Hypotheses', Synthese 37, 260–72.

Kuipers, T. (1984): 'Two Types of Inductive Analogy by Similarity', Erkenntnis 21, 63–87.

Lakoff, G. (1987): Women, Fire, and Dangerous Things: What Categories Reveal about the Mind, Chicago: University of Chicago Press.

Lepore, E. and B. McLaughlin, eds. (1985): Actions and Events: Perspectives on the Philosophy of Donald Davidson, Oxford: Basil Blackwell.

Maher, P. (2000): 'Probabilities for Two Properties', Erkenntnis 52, 63–91.

Margolis, E. and S. Laurence, eds. (1999): Concepts: Core Readings, Cambridge, MA, and London: The MIT Press.

Mill, J. S. (1973–74): A System of Logic, Ratiocinative and Inductive, J. M. Robson (ed.), Toronto: University of Toronto and Routledge & Kegan Paul. Originally published 1872.

Niiniluoto, I. (1980): 'Analogy, Transitivity, and the Confirmation of Theories', in L. J. Cohen and M. Hesse (eds.), Applications of Inductive Logic, Oxford: Oxford University Press, pp. 218–234.

Niiniluoto, I. (1981): 'Analogy and Inductive Logic', Erkenntnis 16, 1–34.

Niiniluoto, I. (1988): 'Analogy and Similarity in Scientific Reasoning', in D. Helman (ed.), Analogical Reasoning, Dordrecht: Kluwer, pp. 271–298.

Ong, W. (1982): Orality and Literacy: The Technologizing of the Word, London: Methuen.

Peirce, C. S. (1902): 'Application for Support for His Logic', Memoir 19, On Arguments, Draft E, http://members.door.net/arisbe/menu/library/bycsp/l75/ver1/m19 (accessed April 22, 2005).

Peirce, C. S. (1931–1958): Collected Papers of Charles Sanders Peirce, ed. C. Hartshorne, P. Weiss, and A. Burks, Cambridge, MA: Harvard University Press.

Pietarinen, J. (1972): Lawlikeness, Analogy, and Inductive Logic, Amsterdam: North-Holland.

Plato (1975): The Trial and Death of Socrates, trans. G. M. A. Grube, Indianapolis: Hackett.

Putnam, H. (1983): Realism and Reason, Vol. 3 of Philosophical Papers, Cambridge: Cambridge University Press.

Reiss, S. (1990): 'The Last Days of Eden', Newsweek, December 3, 44–46.

Rosch, E. (1973): 'Natural Categories', Cognitive Psychology 4, 328–50.

Rosch, E. (1978): 'Principles of Categorization', in E. Rosch and B. Lloyd (eds.), Cognition and Categorization, Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 27–48. Rpt. in Margolis and Laurence (eds.), pp. 189–206.

Rosch, E. (1983): 'Prototype Classification and Logical Classification: The Two Systems', in E. Scholnick (ed.), New Trends in Conceptual Representation: Challenges to Piaget's Theory?, Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 73–86.

Rosch, E. and C. Mervis (1975): 'Family Resemblances: Studies in the Internal Structure of Categories', Cognitive Psychology 7, 573–605.

Rosch, E., C. Mervis, W. Gray, D. Johnson, P. Boyes-Braem (1976): 'Basic Objects in Natural Categories', Cognitive Psychology 8, 382–439.

Schilpp, P., ed. (1963): The Philosophy of Rudolf Carnap, La Salle, IL: Open Court and London: Cambridge University Press.

Seneca (1996): Epistles: 66–92, trans. Richard M. Gummere, Cambridge, MA: Harvard University Press. Ad Lucilium Epistulae Morales originally published 63–65.

Skyrms, B. (1986): Choice and Chance, 3rd ed., Belmont, CA: Wadsworth.

Skyrms, B. (1991): 'Carnapian Inductive Logic for Markov Chains', Erkenntnis 35, 439–460.

Skyrms, B. (1993): 'Analogy by Similarity in Hyper-Carnapian Inductive Logic', in J. Earman, A. Janis, G. Massey, and N. Rescher (eds.), Philosophical Problems of the Internal and External Worlds: Essays on the Philosophy of Adolf Grünbaum, Pittsburgh: University of Pittsburgh Press, pp. 273–282.

Sorabji, R. (1993): Animal Minds and Human Morals: The Origins of the Western Debate, London: Duckworth.

Sorensen, R. (1988): Blindspots, Oxford: Clarendon Press.

Spohn, W. (1981): 'Analogy and Inductive Logic: A Note on Niiniluoto', Erkenntnis 16, 35–52.

Tye, M. (1990): 'Vague Objects', Mind 99, 535–557.

Von Wright, G. H. (1963): The Varieties of Goodness, London: Routledge & Kegan Paul.

Waismann, F. (1930–31): ‘Logische Analyse des Wahrscheinlichkeitsbegriffs’, Erkenntnis 1, 228–248. English trans. (1977): ‘A Logical Analysis of the Concept of Probability’, in F. Waismann, Philosophical Papers, Dordrecht: D. Reidel, pp. 4–21.

Weems, M. (1962): The Life of Washington, Cambridge, MA: The Belknap Press of the Harvard University Press. Weems’ ninth edition originally published 1809.

Welch, J. (1997): 'Analogy in Ethics: Pragmatics and Semantics', in Paul Weingartner, Gerhard Schurz, Georg Dorn (eds.), The Role of Pragmatics in Contemporary Philosophy, Kirchberg am Wechsel: Die Österreichische Ludwig Wittgenstein Gesellschaft, 1016-21.

Welch, J. (1999): 'Singular Analogy and Quantitative Inductive Logics', Theoria 14, 207–247.

Williams, B. (1985): Ethics and the Limits of Philosophy, Cambidge, MA: Harvard University Press.

Williamson, T. (1994): Vagueness, London: Routledge.

Wittgenstein, L. (1922): Tractatus Logico-Philosophicus, London: Routledge & Kegan Paul.

Wittgenstein, L. (1953): Philosophical Investigations, trans. G. E. M. Anscombe, Oxford: Basil Blackwell.